


AD-A111 326

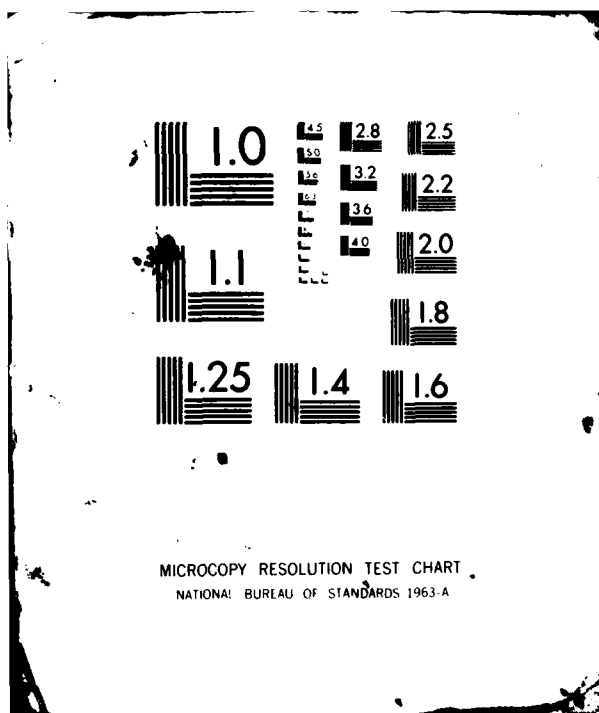
AIR FORCE INST OF TECH WRIGHT-PATTERSON AFB OH F/6 17/9
THE DESIGN AND PERFORMANCE CHARACTERISTICS OF A CELLULAR LOGIC --ETC(U)
APR 81 L A ANKENY
AFIT/DS/EE/81-1


UNCLASSIFIED

NL

1-2
a
b







UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER AFIT/DS/EE/81-1	2. GOVT ACCESSION NO. AD-A111 326	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) The Design and Performance Characteristics of a Cellular Logic 3-D Image Classification Processor		5. TYPE OF REPORT & PERIOD COVERED PhD Dissertation
7. AUTHOR(s) Lawrence A Ankeney Lt Col USAF		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Air Force Institute of Technology (AFIT-EN) Wright Patterson AFB OH 45433		8. CONTRACT OR GRANT NUMBER(s)
11. CONTROLLING OFFICE NAME AND ADDRESS		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS N/A
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE April 1981
		13. NUMBER OF PAGES
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES Approved for public release; IAW AFR 190-17 <i>Fredric C. Lynch</i> FREDRIC C LYNCH 28 JAN 1982 Major USAF Director of Public Affairs		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Pattern Recognition Cellular Automata Cellular Logic Target Classification Neighborhood Transformation Image Processing Laser Radar		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The introduction of high resolution scanning laser radar systems which are capable of collecting range and reflectivity images, is predicted to significantly influence the development of processors capable of performing autonomous target classification tasks. Actively sensed range images are shown to be superior to passively collected infrared images in both image stability and information content. An illustrated tutorial introduces cellular logic (neighborhood) transformations and two and three-dimensional erosion and (continued on reverse)		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

Block 20 (cont'd)

dilation operations which are used for noise filters and geometric shape measurement. A unique "cookbook" approach to selecting a sequence of neighborhood transformations suitable for object measurement is developed and related to false alarm rate and algorithm effectiveness measures. The cookbook design approach is used to develop an algorithm to classify objects based upon their 3-D geometrical features. A Monte Carlo performance analysis is used to demonstrate the utility of the design approach by characterizing the ability of the algorithm to classify randomly positioned three-dimensional objects in the presence of additive noise, scale variations, and other forms of image distortion.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

AFIT/DS/EE/81-1

LEVEL II

①

DTIC
ELECTE
FEB 19 1982
S D E

THE DESIGN AND PERFORMANCE CHARACTERISTICS
OF A CELLULAR LOGIC 3-D IMAGE
CLASSIFICATION PROCESSOR

Doctoral thesis
DISSERTATION

A101143

AFIT/DS/EE/81-1

Lawrence A. Ankeney
LtCol USAF

Approved for public release; distribution unlimited

THE DESIGN AND PERFORMANCE CHARACTERISTICS
OF A CELLULAR LOGIC 3-D IMAGE
CLASSIFICATION PROCESSOR

by

Lawrence A. Ankeney, B.S., M.S.

LtCol

USAF



Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	

Approved:

Date

Matthew Kobusky
Chairman

10 August 1981

Stanley R. Robinson
John J. Line Jr.
Robert W. Carl
John Z. Miller

13 August 1981
21 August 1981
21 August 1981
28 August 1981

Accepted:

JSP Remienicki

28 October 1981

Dean, School of Engineering

AFIT/DS/EE/81-1

THE DESIGN AND PERFORMANCE CHARACTERISTICS
OF A CELLULAR LOGIC 3-D IMAGE
CLASSIFICATION PROCESSOR

DISSERTATION

Presented to the Faculty of the School of Engineering
of the Air Force Institute of Technology

Air University

in Partial Fulfillment of the
Requirements for the Degree of
Doctor of Philosophy

by

Lawrence A. Ankeney, B.S., M.S.

LtCol

USAF

Approved for public release; distribution unlimited

Preface

There have been several trends over the last decade in pattern recognition technology. For example, sensor technology has recently provided a means of collecting high resolution range images (data) of scenes and objects, and many individuals within the pattern recognition community have been quick to recognize the promise of this new form of data. Likewise, several realizations of spatially oriented parallel processors have appeared increasingly promising with respect to providing a real time method of processing the large volume of data which are collected by imaging sensors. Neither of these technological concepts are in reality new or original. However, the recent advances in microelectronics and optical sensors have made the application of these technologies realizable, even within the restrictive cost and size constraints of tactical military systems. Thus, when this dissertation was initiated, the goal was established to investigate analytical methods of predicting the performance of a spatially oriented target classification processor as applied to range images. This goal was reevaluated and revised upon recognizing the following two observations. First, the highly nonlinear behavior of neighborhood transformations generally limit the statistical tractability of cellular logic processors to Monte Carlo performance analysis techniques. Secondly, due to the newness of both image quality range data and cellular logic processor technologies, no generally applicable method of designing 3-D feature measurement algorithms was available. Indeed, only a few individuals were experimenting in this field, and it became obvious that, in general, their approaches to algorithm design more

closely resembled an art form than a science. Therefore, with the sponsorship of the Air Force Avionics Laboratory, the uniqueness of 3-D range data was investigated, and numerous neighborhood transformation algorithm design techniques were developed and evaluated. Eventually, an algorithm design technique emerged which appeared to exhibit the best properties of several earlier approaches. This dissertation tutorially describes this investigation, the algorithm design technique, and some of its performance characteristics.

I would like to sincerely thank my advisor, Dr. Matthew Kabrisky, of the Air Force Institute of Technology for his timely suggestions and encouragement. The critical analysis and support of Dr. John Jones and LtCol Joseph Carl of the Air Force Institute of Technology, Dr. Stanley Robinson and Dr. Stanley Sternberg of the Environmental Institute of Michigan, Dr. Peter Miller of the Perkin Elmer Corp., and Dr. Bradley Sowers of General Dynamics (Convair) are also very much appreciated. I would also like to thank Ms Karen Olin of Hughes Research Laboratories for her helpfulness. Finally, the understanding and patience exhibited by my wife, Marilyn, and daughter, Emily, can neither be underestimated, nor forgotten.

Lawrence A. Ankeney

Contents

Preface.	iii
List of Figures.	vii
List of Tables	ix
Notation	x
Abstract	xiii
I. Introduction.	1
II. A New Breed of Tactical Data.	5
Active vs Passive Data	5
Tactical Target Classification Considerations.	6
Grey Level Interpretation	6
Feature Selection and Modeling.	8
Geometrical Distortions	14
III. Neighborhood Transformations	19
The Basic Transformations.	23
Dilation.	24
Erosion	31
Applications in Shape Measurement	39
Two-Dimensional Closure and Opening.	42
Two Special Purpose Transformations.	52
The Medial Axis Transformation.	52
The End Erosion Operation	53
Three-Dimensional Cellular Logic Operations.	56
IV. The Processor Design.	64
The Design Approach.	64
The Scenario and Assumptions	67
The Targets	67
The Sensor.	69
The Environment	69
The 3-D Noise Model	71
The Processor Architecture	72
Feature Selection.	74
The Noise Filter	76
The 3-D Noise Filter.	77
The 2-D Noise Filter.	85

Feature Measurement.	87
Circle Measurement.	88
Rectangle Measurement	97
The Decision Criteria.	108
Feature Acceptance Window Design.	111
The Target Class Decision	114
V. A Performance Analysis.	119
Performance Defined.	119
The Analysis Procedures and Results.	123
The Basic Processor Performance.	124
False Alarm Rate.	126
Additional Performance Characteristics.	129
Real World Considerations.	134
The Flat Terrain Assumption	134
The Nadir Assumption.	135
The Angular and Range Resolution.	136
Feasibility of Implementation	136
Acceptable Performance.	137
Discussion of Results.	138
VI. Conclusions and Recommendations	141
Bibliography	146
Vita	148

List of Figures

<u>Figures</u>	<u>Page</u>
1 The Diurnal Variations of Passive IR Data.	10
2 The Kirsch Edge Operator Applied to Passive IR Data.	12
3 The Kirsch Edge Operator Applied to Unambiguous Range Data	13
4 Examples of Digitized Three-Dimensional Objects.	16
5 Examples of Neighborhoods.	20
6 Dilation Properties.	26
7 The Many to One Mapping of Dilation.	28
8 A One-Dimensional Example of Dilation.	30
9 The Erosion Operation.	32
10 Erosion Properties	34
11 The Many to One Mapping of Erosion	36
12 A One-Dimensional Example of Erosion	38
13 Equivalent Structuring Elements.	40
14 Examples of Opening.	44
15 Examples of Closure.	47
16 Closure and Opening Geometrical Relationships.	49
17 The Medial Axis Transformation (MAT)	54
18 The End Erode Transformation (EERO).	55
19 Three-Dimensional Closure and Opening.	59
20 The Local Average Function	61
21 3-D Closure and Average Comparison	62
22 Target Class Models.	68

<u>Figure</u>	<u>Page</u>
23 Conceptual Line Scanning Scenario.	70
24 The Processor Architecture	73
25 A Digitized and Thresholded Class 3 Target	75
26 Closure and Opening Applied to a 3-D Image	79
27 Closure and Opening Applied to a Noise Corrupted 3-D Image.	81
28 Closure and Opening Using Structuring Element B	83
29 Typical Thresholded Target Tops.	89
30 Residue Count Relationships.	90
31 Symmetrical Structuring Elements	93
32 Circle Measurement Algorithm Selection	95
33 Conceptual Example of Erosion Sequence Sensitivity to Target Rotation	99
34 Application of the Medial Axis Transformation.	100
35 Typical Residue Set Spread Functions	104
36 End Erosion Shape Discrimination	107
37 Basic Processor Design	110
38 Sample Class Conditional Probability Density Functions (i-th Branch).	112
39 Application of Noise to Signal Ratio	121
40 Basic System Performance	125
41 False Alarm Noise Samples.	128
42 System Performance with Target Scale Variations.	130
43 System Performance with Uncorrected Roll Rate.	132
44 System Performance with Uncorrected Pitch Rate	133

List of Tables

Table	Page
I Preferred Erosion-Mat Sequences for Rectangular Objects.	106
II Branch Acceptance Windows	112
III Strict Decision Criteria.	116
IV Relaxed Decision Criteria	116

Notation

<u>Symbol</u>	<u>Description</u>	<u>Page</u>
\mathcal{Q}	The set of all possible images (events) that may occur within an $I \times J$ matrix	27
\mathcal{Q}_i	The subset of images in \mathcal{Q} that erode to the i^{th} image (E_i) in the range of erosion	35
\mathcal{Q}^i	The subset of images in \mathcal{Q} that dilate to the i^{th} image (D^i) in the range of dilation	29
A	An image (event) in \mathcal{Q}	25
A_i	The i^{th} image (event) in \mathcal{Q}	27
a_{ij}	The ij^{th} element (pixel) of A	25
A^C	The complement of the binary image A	31
A^B	The closure of A with respect to neighborhood B	45
A_B	The opening of A with respect to neighborhood B	45
A_n^{MIN}	The smallest event A in \mathcal{Q} which can erode to the n^{th} event in the range of erosion	37
A_m^{MAX}	The largest event A in \mathcal{Q} which can dilate to the m^{th} event in the range of dilation	31
AW_i	The acceptance window (set of integers) for the i^{th} processor branch	109
B	A set of spatially related cells which defines a neighborhood with respect to a specific (root) cell. Also referred to as a structuring element	19
B'	The reflection of the B neighborhood through its root cell	23
$B(a_{ij})$	The neighborhood (structuring element) B positioned with its root cell superimposed on the ij^{th} cell of the image A	25
B_1	A specific neighborhood (structuring element)	25

<u>Symbol</u>	<u>Description</u>	<u>Page</u>
\mathcal{D}	The set of images (events) in the range of dilation	28
D^m	The m^{th} event in the range of dilation	27
$D_B(A)$	The dilation of the image A with respect to the B neighborhood	25
\mathcal{E}	The set of images (events) in the range of erosion	35
E_n	The n^{th} event in the range of erosion	35
$E_B(A)$	The erosion of the image A with respect to the B neighborhood	31
EERO	The end erosion operation	53
\in	"an element of"	25
$\bar{\Phi}$	A null (all state zero) image	41
H	A noise free image	43
L_i	The logical state (true or false) of the i^{th} branch	114
MAT	The medial axis transform	52
$N(0, \sigma^2)$	Zero mean Gaussian noise with standard deviation σ	71
n	Abbreviation for noise	120
P_d	Probability of detection	65
P_{fa}	Probability of false alarm	65
P_m	Probability of miss	65
P_{mc}	Probability of misclassification	65
r_i	The residue count (number of cells remaining in state "1") at the output of the i^{th} branch	109
s	A value associated with a signal	120
SE $_i$	The i^{th} structuring element	24

<u>Symbol</u>	<u>Description</u>	<u>Page</u>
$T_B(A)$	The transformation of the A matrix with respect to the B neighborhood	19
T_h	A specific threshold value	74
$T_h(A)$	The result of thresholding image A at value T_h	74
\leq	"contained in or equal to"	27
\forall	"for every"	25

Abstract

The introduction of high resolution scanning laser radar systems, which are capable of collecting data in the form of range and reflectivity images, is predicted to have a profound influence on the development of processors capable of performing autonomous target classification tasks. The actively sensed range images are shown to be superior to passively collected infrared images in two areas: the actively sensed range images are relatively insensitive to diurnal and environmental variations, and the range images provide a direct measure of the scene's three-dimensional shape rather than depending on thermal variations to provide shape information. Additionally, the amount of laser energy reflected from each surface in the scene can be measured, thus providing an additional reflectivity image of the scene. An approach to processing range images via cellular logic (neighborhood) transformations is described, and a unique neighborhood transformation algorithm selection procedure is developed and generalized to an extent that geometric shape measurement algorithm selection can be performed in a "cookbook" fashion. The concepts of residue set and residue set spread functions are defined and shown to provide significant insight to the feature extraction effectiveness and potential false alarm rate of the algorithm under consideration. A hypothetical scenario and processor architecture are described and the algorithm design approach is used to select a sequence of neighborhood transformations which perform three-dimensional feature measurement of rectangular box and truncated cone shaped geometrical objects. A Monte Carlo performance analysis is used to demonstrate the utility of the design approach by characterizing the ability of the processor to classify randomly

positioned three-dimensional objects in the presence of additive noise, scale variations, and other forms of image distortion. An illustrated tutorial is provided to introduce the concept of neighborhood transformations and to develop the two and three-dimensional erosion and dilation operations which are used for noise filtering and feature measurement within the proposed processor design.

I. Introduction

For many years man has attempted to construct machines which perform tasks similar to those easily accomplished by most humans. For example, machines capable of uniquely classifying two-dimensional geometric objects such as printed letters, finger prints, or audiographs have yet to be designed for use in an environment where variables such as scale, rotation, font, background, energy, and noise cannot be controlled. The problems associated with the classification of three-dimensional objects compound these difficulties by adding a third spatial dimension within which the objects can be translated and rotated. Thus, very little success has been demonstrated in designing machines which can autonomously (i.e. without human intervention) acquire and classify tactical vehicles such as trucks, tanks or jeeps in an uncontrolled environment such as a battlefield. Nevertheless, the desirability of obtaining such a machine is obvious.

The difficulties encountered in attempting to totally describe tactical targets and their backgrounds in either a Gestalt or mathematical sense can generally be attributed to either the diurnal variability of the data (which prohibits modeling the objects or background as a stationary random process, Ref. 23) or the inability to measure scene information that is directly related to the geometrical shape of objects within the scene. Passively collected infrared (IR) data is a classic example of a data base which exhibits both these shortcomings. Since passive IR data is a measure of the thermal emissivity of the scene, it is very sensitive to diurnal variations such as time of day, solar loading, moisture content and, to a certain

extent, environmental history. Likewise, thermal emissivity offers only an indirect and relative measure of an object's shape. Thus, designers of target classification algorithms have been forced to use ad hoc features and suboptimal feature extraction/measurement techniques. A more subtle but extremely relevant issue associated with the inability to model targets, noise, and clutter is that, without such a model, volumes of real world data must be collected for feature selection, algorithm training, and processor performance evaluation purposes. Additionally, more often than not, sensor peculiarities tend to make the data collected with one sensor essentially useless for designing or evaluating target classification processors which will use data collected by other sensors. As a result, the expense associated with collecting a statistically significant data base is inevitably prohibitive and the data available for feature selection, algorithm training, and evaluation is generally limited.

Fortunately, recent advances in sensor technology and practical realization of parallel processor architectures may give a new lease on life to the tactical target classification community. The Air Force Avionics Laboratory has, through the Environmental Institute of Michigan, recently demonstrated a limited capability to classify (and possibly identify within classes) tactical targets in real time by processing high resolution image quality range data in a spatially organized parallel processor not dissimilar in design to that proposed by Unger (Ref. 25). Even though other commercial and academic institutions have also indicated varying degrees of success in this area, the algorithm design and selection process of such a machine has, unfortunately, remained an art and performance evaluations have

generally been limited to a few samples of real world data. This dissertation will attempt to remove the mystery surrounding the uniqueness of actively collected range data and to exploit the consistency of this format of data by developing a unique "cookbook" approach to processor algorithm selection and to evaluate the performance trends of a processor designed in accordance with this cookbook. Chapter II will first describe the virtues and uniqueness of actively collected data as compared to the various forms of passive data which have been historically available to the tactical target classification community. Chapter III will then briefly describe cellular logic processors and describe in greater detail a useful variety of cellular logic operations (also known as neighborhood transformations) and provide a tutorial description of some of their applications and properties. Chapter IV then develops a step by step approach to the design of a target classification processor which uses cellular logic operations. The emphasis of Chapter IV is in using the three-dimensional information content of the data and providing an approach to shape classification algorithm selection which, when combined with an appropriate decision criteria, supports low probability of error and false alarm rate performance goals. Chapter V provides an extensive Monte Carlo performance and parametric sensitivity analysis of this processor and attempts to relate the demonstrated performance to the operational appropriateness of the assumptions and design approach. Chapter VI then summarizes the conclusions and provides recommendations for future studies.

While a specific set of hypothetical targets is proposed, features selected and extracted, and performance indicated for a given processor

architecture, the reader's emphasis should remain with understanding the advantages of working with active imagery and the various tradeoffs associated with the proposed cellular logic design approach. For, without understanding these issues, these technologies will remain an art form.

II. A New Breed of Tactical Data

Historically, two generic forms of data have been available to the tactical target identification community: actively collected data where energy is transmitted and some measure of the reflected signal is detected, and passively collected data where the detector senses only naturally occurring reflected or radiating information. This chapter will first briefly discuss the qualitative information content of each data type and then compare these to actively collected image quality range data. These comparisons will be made in the context of the specific issues or problem areas associated with the tactical target classification task.

Active vs Passive Data

Active sensors are generally capable of measuring the intensity of the reflected signal as well as the phase or time delay between the transmitted and received pulses. The intensity (energy) of the return provides a relative measure of the surface reflectivity and the phase can provide either an absolute or relative measure of range. Antenna size and processing limitations imposed on tactical missile airframes have not permitted active radar sensors sufficient angular resolution to effectively extract three-dimensional target shape information from the scene. Passively collected data such as television, photography and infrared (IR) are capable of demonstrating image quality resolution. However, range information is not readily accessible in passively sensed data, and a direct measurement of a scene's three-dimensional information content is not available. Thus, any information contained in passively collected data relative to the scene's three-dimensional

shape must be extracted indirectly from the multigrey level intensity data. An ideal tactical sensor might well combine the high resolution capabilities of an imaging optical or IR sensor with the accurate range measuring capabilities of an active sensor to provide image quality three-dimensional range data.

The remainder of this chapter will describe specific tasks or problems associated with tactical target identification and while doing so illustrate the potential improvements (relative to passive data) that three-dimensional range data provides. As a convenient notation, the use of the term 3-D data will henceforth refer to multigrey level image data in which the third dimension grey level provides relative or absolute range information and the term 2-D data will refer to multigrey level image data in which the third dimension grey levels represent scene information other than range or height. Note that 2-D binary data may refer to either thresholded (clipped) 3-D or 2-D data.

Tactical Target Classification Considerations

Every shape or object classification task has its unique and complex idiosyncrasies and, without the existence of a general theory which can be applied to all pattern classification tasks, one must address the "physics" associated with the particular task at hand. In such a manner, this section will briefly describe three specific data related issues which pertain specifically to the task of identifying targets in a tactical military environment. Relevant differences between passively collected 2-D data and actively collected 3-D data will be emphasized.

Gray Level Interpretation. As described earlier, passively collected data depends solely upon sensing information which is inherent

in the scene such as relative thermal temperature or reflection of natural radiation. As a result, passively collected data is sensitive to local temporal and spatial variations in the environment such as ambient temperature changes, precipitation, humidity and solar radiance. Hence, the grey levels of a given scene will in general not be ergodic and will be statistically nonstationary (Ref. 23). Figures 1(a) and (b) provide an example of the variable nature of passively collected infrared data. There was a twelve hour delay between the collection of the first and second images. Note that not only have the relative grey levels of the scene changed, but the grey level rates of change (the intensity gradients) have also varied noticeably. Variations such as these are referred to as diurnal variations and are common to all forms of passively collected imagery. To compound the problems associated with diurnally variant data, passively collected data contains no direct measurement of the scene's three-dimensional shapes or volumes. As a result, passive data requires the use of features such as edge gradients, corners, and areas of constant texture (which may or may not be relevant to the true three-dimensional content of the scene) for the target classification task.

Actively collected data, on the other hand, provides significant improvements in the ability to interpret the grey level information. First, because the sensor has its own radiating source and the velocity of that energy is essentially constant, the reflected return signal will provide (within the design limits of the sensor) the same time delay or phase shift regardless of most environmental variations. Of course there will be exceptions such as the loss of foliage from deciduous trees or shrubs and the accumulation and drifting of snow, but in

general, actively collected 3-D data is relatively insensitive to the diurnal variations which have historically plagued passively sensed data. Secondly, no longer must a system designer attempt to indirectly extract three-dimensional information from the imagery. The imagery is the three-dimensional information! In this context, actively collected 3-D data is superior to passively collected 2-D data.

Feature Selection and Modeling Considerations. The selection of features or discriminants which permit the identification of targets and the rejection of nontarget objects and noise are of primary importance in all pattern recognition tasks. There are generally an unlimited number of features or combinations of features which may be used to separate tanks from trucks or trees such as weight, mobility, color, smell or audible emissions. However, for the task of classifying objects on the ground as viewed from an airborne platform, the most intuitively obvious features should be directly related to the three-dimensional shape of the target.

As discussed in the previous section, the diurnal variability and lack of directly measurable three-dimensional shape information in passively collected data has forced target classification algorithm designers to use various gray level normalization techniques and gradient features such as edges, corners and texture. Since much effort has been expended in developing and characterizing these techniques, there is no need to abandon their use unnecessarily. However, before applying these techniques, their application to 3-D data must be reviewed and understood. Edges, for example, are no longer gradients associated with changes in the scene's reflectivity, color or temperature (as in 2-D data) but now become the points of two

intersecting surfaces. Likewise, areas of constant texture represent surfaces in 3-D data and corners indicate the intersection of three surfaces. Additional features now become available such as volume measurements and height to length or width ratios and a high pass filter becomes a potential terrain removal operation under certain scenarios such as 3-D data collected from a look-down perspective. Thus, 3-D data does not necessarily require a new set of features, noise filters, or normalization techniques, but whatever techniques are used must be understood as applied to the 3-D range information.

As a means to illustrate the results of applying a common operation to both 2-D and 3-D data (and to demonstrate the insensitivity of 3-D data to diurnal variations), an edge detection operation will be applied to the images of Fig. 1(a) and (b) as well as to the same scene as viewed with 3-D actively collected range data. The gradient measurement technique selected for this example is the well known Kirsch operator (Ref. 2) which is defined at the point X as

$$\max_{i=0}^7 [1, \max[(a_i + a_{i+1} + a_{i+2}) - 3(a_{i+3} + \dots + a_{i+7})]] \quad (1)$$

where the subscripts are evaluated modulo 8 for the neighborhood

$$\begin{array}{ccc} a_0 & a_1 & a_2 \\ a_7 & X & a_3 \\ a_6 & a_5 & a_4 \end{array}$$

of X. If a threshold operation is applied to the image following the gradient operator, edges can be defined as those gradients which have a magnitude exceeding the threshold value. This threshold may either be adaptive (related to the magnitude or variability of the gradients) or a constant value depending on the complexity allowed the processor. The Kirsch operator of Eq. 1 was first applied to the passively collected (infrared) intensity images of Fig. 1, and an edge threshold was then




Fig. 1. The Diurnal Variations of Passive Infrared (IR) Data

applied to obtain the edge images of Fig. 2(a) and (b). Note that the edges in these images correspond to the intensity transitions (i.e. changes in surface temperature) in the scenes of Fig. 1(a) and (b) respectively. Also note the obvious differences between the edge images of Fig. 2(a) and (b). Since the same threshold was applied to both gradient images, these differences can be attributed almost exclusively to the diurnal variations exhibited by the passively collected data during the 12 hour delay between scene samples. To contrast the edge images of the passively collected scene data, Fig. 3(a) and (b) provide the comparable edge images for the same scene but sampled with an active (laser) range measuring sensor. Note in this case that the edges are no longer related to the surface temperature of the scenes but instead correspond to discontinuities in the range data. Also, the edge images are almost identical which clearly illustrates the insensitivity of the actively collected range data to diurnal variations in the scene.

As a final observation concerning the selection of features, when range (or height) imagery is used, it may not be necessary to collect and analyze large quantities of imagery data (as is necessary for passive IR data) in the search of target features which are invariant to diurnal variations. The diurnal consistency of 3-D data, when combined with the true three-dimensional shape information contained in the data, permits the a priori selection of potential target features (before seeing any data). And, for the first time, tactical targets, clutter, and terrain can be realistically modeled (either in software or physical scale models) for Monte Carlo simulations or statistically modeled for analytical performance comparisons. Conceptually, similar modeling techniques could be applied to passively collected sensor information,




Fig. 2. The Kirsch Edge Operator Applied to Passive IR Data




Fig. 3. The Kirsch Edge Operator Applied to Range Data

C but the necessity to be able to model many possible environmental, seasonal and topographical scenarios not only makes the model quite complex but also requires an extensive data base against which to design the model. Such target and clutter modeling attempts have been and are presently being made for passive IR scenes (Ref. 23), but success has been limited primarily by the lack of a sufficiently robust data base for design and validation purposes. In comparison, the generation of a complete three-dimensional target and clutter model should be a low risk task requiring only limited amounts of 3-D data for training and validation.

C Geometrical Distortions. Just as a sample and hold circuit provides a less than perfect discrete representation of a continuous time varying wave form, the discrete range sampling of a three dimensional scene also will result in a less than perfect discrete representation of the scene. The spatial distortions which result from such a digitization process are referred to as digitization or quantization noise and the specific digital representation of a scene is in general quite sensitive to the relative position of the scene with respect to the digitizing grid. Additionally, in a system that requires a finite amount of time to sample the scene, any motion of the digitizing grid relative to the scene or objects in the scene will induce additional distortions in the form of twisting or stretching of the discrete image. In a tactical environment, the unpredictable motion, translation, and rotation of objects of interest in the scene, the dynamic capabilities of tactical airborne sensor (data collection) platforms, and the finite spatial resolution of the sensors themselves, combine to provide sufficient justification to discuss these forms of

geometrical distortions in detail.

Figure 4(a) provides an example of the spatial variations which occur when the height of a continuous three-dimensional hatbox shaped object is placed upon and discretely sampled by a regular digitizing point matrix. (Such a digitizing concept is not dissimilar to the capturing of regularly positioned spikes by a randomly thrown hoop (Ref. 12) so often referred to in geometrical probability discussions.) It can be seen in Fig. 4(a) that when moved to a different position on the digitizing matrix, the digitized representation of the hatbox exhibits a quite different shape. Note that once the continuous hatbox object is positioned on the digitizing matrix, it can be rotated around its central axis without affecting its digitized representation. Objects which exhibit this property are defined to be rotationally invariant. Figure 4(b) provides the results of digitizing two samples of a shoebox shaped object which is sensitive to both rotation and translation. In this case, the digitized image of the rotated and translated object has had one of its corners "rounded off" and its width and length are no longer constant. Quite often the geometrical distortions associated with digitization are surprising, especially in small (with respect to the resolution of the digitizing matrix) objects.

Not only will static translation or rotation of a scene (with respect to the digitizing grid) induce various forms of geometrical distortions, any dynamic motion between the scene and the digitizing matrix will also induce additional geometrical distortions to the digitized image. For example, if the digitizing matrix sampled height sequentially from left to right, one row at a time, and the hatbox and shoebox objects described earlier were moving at a constant velocity




Fig. 4. Examples of Digitized Three-Dimensional Objects

across the digitizing matrix, the resulting digital representations of these objects would be skewed or stretched.

An equivalent method of portraying both the static and dynamic digitization processes described above can be realized by positioning the continuous objects (targets) on a surface (terrain) and sampling the range from an airborne platform to the surface with a line scanning range sensor. If the sensor platform passes over the terrain at a velocity so that each scanned line (row) of data neither over nor under samples the previous or next line of information, a topographical representation of the digitized objects and terrain would be obtained. Likewise, in this portrayal, geometrical distortions such as skewing or stretching would be introduced by either movement of the objects beneath the sensor or a constant roll or pitch rate in the sensor platform. Finally, since the resolution cell size is proportional to the altitude at which the sensor is carried above the terrain, a change in sensor altitude will result in scale changes to the digitized image.

Unfortunately, geometrical distortions such as those described above (digitization noise and sensor/target motions) are pervasive in both active and passive data collection efforts. In systems which use passive data, however, digitization noise and small platform instabilities have not generally been of great concern to target acquisition algorithm designers because the diurnal variability of the data has demanded the lion's share of design effort. In systems designed to use 3-D data, diurnal variations should no longer be the driving limitation to algorithm development and, if the designer so desires, the modeling of the above described geometrical distortions for inclusion in performance analyses can be justified by the true

three-dimensional information contained in the 3-D data.

In summary, the most obvious and potentially far reaching advantages that image quality actively collected range data has when compared to similar resolution passively collected data would be the ability to directly measure information which describes the topographical (three-dimensional shape) content of the scene and that this measure of information is not sensitive to most diurnal variations of the environment. Related to these issues are potentially significant improvements in the understanding and selection of features which are directly related to the three-dimensional shape of the target, the potential to select in an a priori manner target features, and the ability to construct realistic target and clutter models for digital system performance analysis. Chapters IV and V will consider each of these characteristics in the design of the target classification processor and its performance analysis. However, Chapter III will first describe the spatial operations known as neighborhood transformations. A clear understanding of these cellular logic functions (and their structuring element equivalent operations) is needed to understand the design of the shape classification processor of Chapter IV.

III. Neighborhood Transformations

Neighborhood transformations (also referred to as cellular logic operations) refer to a generic class of spatially oriented operations which are closely related to the studies of computational geometry (Ref. 16) and cellular automata (Ref. 3). In a macroscopic sense, a neighborhood transformation operates on an entire array of cells to create a transformed array. The transformation, however, is a local operator since the state of each cell in the new transformed array is only a function of its present state and the states of its neighboring cells. More specifically, the neighborhood transformation (T) of the array (A) is defined as

$$T_B(A) = T(A, B) \quad (2)$$

where $T_B(A)$ is the transformed array and B is a local neighborhood of cells which can be arbitrarily specified. In general, the five cell von Neumann and the nine cell Moore neighborhoods of Fig. 5 are often used, and the functional T can be any algebraic or Boolean expression. The Kirsch operator described in Chapter II is an example of a cellular logic operation. Before additional transformations are presented, a brief review of pertinent historical and technical issues relating to cellular logic operations and their mechanization is appropriate.

The early works of von Neumann (Ref. 26) were centered around theoretical cellular automata concepts such as machine self-reproduction. His work is described and the early theoretical works of other authors are compiled in a comprehensive set of essays edited by Burks (Ref. 3). A summary of more recent theoretical efforts is




Fig. 5. Examples of Neighborhoods

provided in the IEEE Proceedings on Cellular Automata, and the work of Rosenfeld and Dyer (Ref. 22) at the University of Maryland illustrates one of the more unique constructs being investigated.

The mechanization of a two-dimensional spatially organized parallel processor was described by Unger (Ref. 25) in 1958. Since then, numerous special purpose parallel processors have been designed to perform various pattern recognition and image processing functions. While the University of Illinois ILLIAC III (Ref. 14) and the CLIP (Ref. 6) machines designed at University College London appear to be the most computationally powerful (Ref. 20), the Perkin-Elmer Corp. CELLSCAN (Ref. 10) and GLOPR (Ref. 21) processors have enjoyed some success in the commercial marketplace. The Environmental Institute of Michigan (ERIM) is presently under contract with Harris Semiconductor through the Air Force Avionics Laboratory to construct a large scale integrated (LSI) circuit design of a programmable cellular logic stage which promises to improve real time image processing capability. Throughout this period of hardware development, many articles have been published which describe useful geometrical and topological properties of cellular logic operations. For example, McCormic (Ref. 14) and Golay (Ref. 7) established a foundation for both regular hexagonal and rectangular tessellation organizations and Grey (Ref. 8) addressed cellular connectivity, Euler number measurement, and perimeter estimates. Several French authors have also been active in this field, and their approach of describing neighborhood transformations as "hit or miss" operators is unique as is their concept of structuring element operations. Some of their more descriptive efforts have been translated and published in the Journal of Microscopy (Ref. 11). In summary, the

field of cellular logic processing has had a rich theoretical basis, and a diverse number of investigators have actively participated in applying cellular logic operations to image processing and pattern recognition tasks. For a more detailed description of the theoretical aspects, the bibliography by Nishio (Ref. 18) and article by Maruoka (Ref. 13) are most recent. On the other hand, Preston, Duff, Levialdi, Norgren, and Toriwaki (Ref. 20) provide a very complete summary of the more applications oriented efforts in this field and include a very complete (with the exception of the French authors mentioned above) list of references.

The remainder of this chapter will describe specific neighborhood transformations and sequences of transformations which will be used later in the design of a target classification processor. Emphasis is placed on developing and interpreting the geometrical relationships associated with applying these operations to binary images in two and three dimensions. The original tutorial instruction provided and illustrative examples (which permit the logical transition of two-dimensional operations to binary three-dimensional operations) will be most appreciated by anyone who has attempted to extract and combine similar information from any of the above referenced sources.

First to be described will be the dilation and erosion operations which are the basic building block transformations. These two operations are also referred to as expand and shrink functions, and the specific geometrical operations they perform are often in agreement with their descriptive names. Dilation and erosion have been used by Miller (Ref. 15) to define a Boolean algebra which describes many cellular logic functions. The closure and opening operations are then described.

These cellular logic operators are constructed by simple sequences of dilation and erosion transformations, and they perform interesting geometrical bounding functions. While neighborhoods such as the von Neuman and Moore are commonly used, certain less common neighborhoods will be described which are useful for more complex transformations which skeletonize images and erode ends of line segments. Finally, the concept of binary cellular operations will be expanded to three-dimensions (multigrey level data) where dilations, erosions, closures and openings are also quite useful. An additional sequence of three-dimensional transformations which demonstrate properties similar to a high pass filter when applied to three-dimensional surfaces is described by Sternberg (Ref. 24) but is not used in the design of the processor of Chapter IV. Suggested uses for these transformations are discussed throughout the chapter, and emphasis is placed on describing the capabilities and limitations of each in a tutorial manner. Specific designs and applications will not be described until Chapter IV.

The Basic Transformations

Before defining the dilation and erosion transformations, terminology common to neighborhood transformations will be reviewed. The concept of a neighborhood simply refers to one or more spatially related cells. However, a neighborhood (of cells) is always specified with respect to a root cell. Figure 5 provides several examples of neighborhoods, each with its root cell marked with a "+". A symmetrical neighborhood is any neighborhood which remains unchanged when reflected through its root cell. Therefore, neighborhoods B_1 and B_2 of Fig. 5 are unsymmetrical (note that B_1 is the reflection of B_2 through its root cell) and the remaining neighborhoods are symmetrical with respect to

the specified root cells.

Quite often, neighborhood transformations are designed to perform simple Boolean operations on binary images or binary surfaces (a binary surface is multigrey level data which can be partitioned so that the volume of cells on and below the surface are in state "1" and the volume of cells above the surface are in state "0"). In these cases, a specified neighborhood can be considered a geometrical shape constructed of state "1" cells (i.e. a structuring element) and the desired transformation can be effected by simple binary additions and products as the structuring element is appropriately positioned throughout the image. Neighborhood transformations will be defined in terms of structuring element operations whenever possible because they are easier to describe (both conceptually and mathematically) and, as will be seen shortly, the structuring element concept permits a direct means of tracking the geometrical interpretations associated with sequences of dilations and erosions. The concepts of root cells and symmetry as described for neighborhoods apply also to structuring elements. The basic concepts of dilation and erosion will now be defined in terms of structuring element operations. Since the relationships described in the following sections have either been proven by Miller (Ref. 15) or demonstrated by Matheron et.al. (Ref. 11), their definitions and descriptions are graphically illustrated and comparatively discussed to provide the reader an understanding of the operations and their potential applications.

Dilation. The dilation operation as applied to the state "1" cells of a binary image (matrix)

$$A = \{a_{i,j}\} \quad (3)$$

is defined as

$$D_B(A) = \sum B(a_{i,j}) \forall a_{i,j} = 1, a_{i,j} \in A \quad (4)$$

where $B(a_{i,j})$ is a structuring element (neighborhood of state "1" cells) positioned with its root cell at the $a_{i,j}$ cell of the image and the summation specifies binary addition. Stated geometrically, the dilation operation specifies the set of cells spatially covered by the structuring element $B(a_{i,j})$ as it is positioned at each state "1" cell of the binary image A . Figure 6(a) illustrates a binary image (A) and two structuring elements B_1 and B_2 . The dilation of the image (A) using the structuring elements B_1 and B_2 is illustrated in Fig. 6(b) as $D_{B_1}(A)$ and $D_{B_2}(A)$ respectively. Note that the dilated images consist of the original image (crosshatched in Fig. 6(b)) and those cells covered by the respective structuring elements state "1" cells as the structuring elements root cell is positioned at each state "1" cell in the original undilated image.

The dilation operation has commutative and associative properties (Ref. 15) such that

$$D_{B_1}(D_{B_2}(A)) = D_{B_2}(D_{B_1}(A)) \quad (5)$$

and

$$D_{B_1}(D_{B_2}(A)) = D_{B_3}(A) \quad (6)$$

where




Fig. 6. Dilation Properties

$$B_3 = D_{B_1}(B_2) = D_{B_2}(B_1) . \quad (7)$$

Thus, a sequence of dilations can be implemented by a single operation using the larger structuring element B_3 of Eq. 7. Figure 6(c) illustrates the sequence of dilating the image (A) first with B_1 and then with B_2 and the resulting image is labeled $D_{B_2}(D_{B_1}(A))$. Figure 6(d) then illustrates the structuring element B_3 and the dilation of the image (A) using this larger structuring element. The verification of Eq. 5 proceeds in the same manner. The definition of dilation insures that the dilated image will always be at least as large as the original image so that

$$A \leq D_B(A) \quad (8)$$

where the notation " \leq " is read "contained in" and implies that the spatially positioned set of state "1" cells in the image (A) is a spatial subset of the state "1" cells of $D_B(A)$ so that the binary product

$$A \cdot D_B(A) = A \quad (9)$$

holds true. The equality of Eq. 8 holds for the structuring element which consists of only a root cell.

Dilation is a many to one mapping. Figure 7(a) conceptually illustrates the domain and range relationship for dilation, and Fig. 7(b) provides an example of three images (A_1 , A_2 and A_3), each of which dilates to the same image D^X . The domain (\mathcal{Q}) of dilation can be bounded by specifying that it consist of the set of all possible events (binary images) that may occur on an N by N segment of a larger M by M




Fig. 7. The many To One mapping of Dilation

matrix. It is also assumed that all the cells outside the N by N matrix are in state "0" and that M is sufficiently larger than N so that any dilations applied to an image in the domain would not create an image in the range that could not be contained in the M by M matrix. Since each event of the range is the image of one or more events in the domain, the dilation operation is onto (surjective) and the domain can be partitioned into sets of events, each member of which, dilates to a specific event in the range:

$$\mathcal{Q}^m = \{A : A \in \mathcal{Q}, D_B(A) = D^m, D^m \in \mathcal{D}\}. \quad (10)$$

Figure 8 provides a one-dimensional example of the domain and range relationships associated with a four cell ($N=4$) portion of a larger ($M>4$) bit stream. The 2^4 possible events in the domain (\mathcal{Q}) dilate to one of five events in the range (\mathcal{D}) when the one-dimensional structuring element B is used. Note that some, but not all, of the events $D \in \mathcal{D}$ are also in \mathcal{Q} and visa-versa. For example, the null event in the domain (A_0) maps to the null event in the range (D^0), the events A_1 thru A_4 map (allowing translation) to the three cell event D^1 , and A_7 maps to the four cell event D^2 . As an example of a set \mathcal{Q}^m , the events A_{12} , A_{13} , A_{14} , and A_{15} of Fig. 8 each dilate to the event D^4 of the range, thus

$$\mathcal{Q}^4 = \{A_{12}, A_{13}, A_{14}, A_{15}\}. \quad (11)$$

It is obvious from the definition of dilation that the set of events (\mathcal{Q}^m) will contain an image




Fig. 8. A One-Dimensional Example of Dilation

$$A_{MAX}^m = \text{MAX}[A \in Q^m] \quad (12)$$

which is as large as or larger (contains more state "1" cells) than the other members of the set. Event A_{11} of Fig. 8 is an example of this concept because, if another state "1" cell were added to A_{11} (or A_{10}), it would dilate to an image larger than D^3 . Thus, the event A_{MAX}^m is an upper bound on the size of events that can dilate to the event D^m . Image A_3 of Fig. 7(b) is another example of an image which is the largest event that can dilate to the event D^x using the Moore structuring element. It will be described later in this chapter that the event $A_{MAX}^m \in Q^m$ can be determined by applying an erosion operation to the event $D^m \in \mathcal{D}$. This and other relationships between dilation and erosion as well as potential uses for dilation in image processing will be explored in the following sections.

Erosion. The second commonly used neighborhood transformation is the erosion operation. Erosion is defined in terms of the dilation operation as

$$E_B(A) = (D_B(A^c))^c \quad (13)$$

where A^c is the complement of the binary image A , and it is formed by changing the state "1" cells in the image (A) to state "0" and state "0" cells to state "1". Figure 9 illustrates a step by step implementation of Eq. 13 as applied to the binary image labeled A and the structuring element labeled B . The complement of the image A is first calculated and then A^c is dilated using the unsymmetrical structuring element B to form the image labeled $D_B(A^c)$. The erosion operation is then completed




Fig. 9. The Erosion Operation

by taking the complement of $D_B(A^c)$ and the result is labeled $E_B(A)$.

An equivalent (Ref. 11) erosion structuring element operation is

$$E_B(A) = \{a_{i,j} : B'(a_{i,j}) \leq A, a_{i,j} = 1, a_{i,j} \in A\}. \quad (14)$$

Using this definition of erosion, one can readily verify that the reflected structuring element (labeled B') in Fig. 9 can only be contained in the image A at the same three positions earlier determined to be the erosion of A . Stated more graphically, the erosion $E_B(A)$ is the set of positions at which the structuring element B' fits into the state "1" cells of the binary image A . Figure 10(a) illustrates a binary image labeled A and two structuring elements labeled B_1 and B_2 . The erosion of the binary image (A) using B_1 and B_2 are illustrated in Fig. 10(b) as $E_{B_1}(A)$ and $E_{B_2}(A)$ respectively. Note that these eroded images consist only of the root cell positions at which the reflected structuring element could be contained entirely within the state "1" cells of the image A .

Sequences of erosions do not generally exhibit the same properties (Ref. 15) as dilations. However, as in dilations, the order in which a sequence of erosions is performed does not influence the results of a given erosion sequence applied to an image. Thus,

$$E_{B_1}(E_{B_2}(A)) = E_{B_2}(E_{B_1}(A)). \quad (15)$$

And, as with dilation, a single erosion can replace a sequence of erosions such as

$$E_{B_1}(E_{B_2}(A)) = E_{B_3}(A) \quad (16)$$




Fig. 10. Erosion Properties

where

$$B_3 = D_{B_1}(B_2) = D_{B_2}(B_1). \quad (17)$$

For example, Fig. 10(c) illustrates the sequence of eroding the image A of Fig. 10(a) with structuring elements B_1 and B_2 sequentially, and the resulting image is labeled $E_{B_2 B_1}(A)$. Figure 10(d) then illustrates Eqs. 16 and 17 by first constructing the structuring element B_3 by dilating B_2 with B_1 (or equivalently dilating B_1 with B_2) and then eroding the image A using this equivalent structuring element (B_3). Note that the reflected equivalent structuring element (B'_3) fits in the image (A) at only one position and that this position is the same cell that remained following the sequence of erosions in Fig. 10(c). One might also observe that the definition of erosion insures that the eroded image will always be no larger than the uneroded image. Thus,

$$E_B(A) \leq A, \quad (18)$$

and the equality holds for the structuring element which consists of only one cell, the root cell.

Erosion, as was dilation is also a many to one function. Figure 11(a) conceptually illustrates the domain (\mathcal{Q}) and range (\mathcal{E}) relationships for erosion and Fig. 11(b) provides an example of three images (A_1, A_2, A_3) each of which erodes to the image labeled E_X . If the domain of erosion is defined as the set of all possible binary events which can occur on a bounded matrix, then each event of the range is the image of at least one event in the domain (the erosion operation is onto), and the domain can be partitioned into sets of events that erode to common events in the range:




Fig. 11. The Many To One Mapping of Erosion

$$Q_n = \{A : A \in Q, E_B(A) = E_n, E_n \in \xi\}. \quad (19)$$

Figure 12 again illustrates the 16 events which may occur in the domain of a one-dimensional four cell matrix and demonstrates the erosion of these events using the structuring element B. In this case, the range consists of only three events. The null event E_0 is the image of the events in the set

$$Q_0 = \{A_i\}, i = 0,1,2,\dots,8,9,12,13,14, \quad (20)$$

E_1 is the image of A_{10} and A_{11} , and E_2 is the image of event A_{15} .

Referring once more to the definition of erosion (Eq. 14), observe that a minimum sized image in the domain must exist for each event in the range. For a given structuring element (B) and an event in the range ($E_n \in \xi$), this minimum event is specified by the set of cells covered by the reflected structuring element as its root cell is positioned at each state "1" cell in the range event E_n . Obviously, any image which does not contain this minimum event cannot possibly erode to event E_n using the structuring element B. Thus, this minimum event

$$A_n^{MIN} = \sum B'(a_{i,j}) \forall (a_{i,j} \in E_n, a_{i,j} = 1). \quad (21)$$

must be contained in each event $A \in Q_n$. Stated in terms of dilation (note the similarity between Eq. 21 and Eq. 4),

$$A_n^{MIN} = D_{B'}(E_n). \quad (22)$$

For example, the event A_3 of Fig. 11(b) is the smallest event which can erode to the event labeled E_x and that the dilation of E_x using the




Fig. 12. A One-Dimensional Example of Erosion

structuring element B' results in the image A . Finally, note that the event A_{MAX}^m of Eq. 12 can be specified by

$$A_{MAX}^m = E_B(D^m) \quad (23)$$

for an event D^m in the range of dilation and structuring element B . The relationships of Eqs. 22 and 23 are basic to understanding the closure and opening operations which are described later in this chapter. However, as an aid to better understanding the dilation and erosion operations, a brief discussion concerning their use in measuring geometrical shapes will first be provided.

Applications in Shape Measurement. A sequence of one or more dilations will generally not be useful in shape discrimination unless applied in conjunction with one or more erosions. One exception to this observation is that dilations may be used to digitally construct geometrical shapes which can then be used as decision templates. Erosions, on the other hand, can be directly used to measure the geometrical size and shape of binary images.

Consider the image A of Fig. 13(a) and assume that our goal is to measure (i.e. geometrically describe in some sense) this image. Using the structuring elements B_1 and B_2 of Fig. 13(a), a sequence of two erosions using structuring element B_1 followed by a single erosion using structuring element B_2 erodes image A to the four cell image labeled $E_{B_x}(A)$. Recalling the definition of erosion (Eq. 14) and its associative properties (Eq. 16 and 17), the image $E_{B_x}(A)$ describes the four positions at which the equivalent structuring element (B'_x) of Fig. 13(b) fits into the image A . (Note that for symmetrical structuring elements, $B = B'$.) If, instead, the image A were eroded twice using structuring elements B_1




Fig. 13. Equivalent Structuring Elements

and twice again using B_2 , the resulting image would be the null set (no state "1" cells) because the equivalent structuring element for this sequence of erosions, labeled B_y in Fig. 13(b), cannot fit into the image A at any position. Thus, one might conclude that the structuring element B_x was smaller than the image A and that B_x could fit into the image A at four positions but that the structuring element B_y was too large to fit into the same image. In other words, we have determined that the image A was larger than the structuring element B_x and smaller than the structuring element B_y . Unfortunately, this knowledge does not totally characterize the image.

It should be quite apparent that the use of erosion sequences to extract shape information from an image is highly dependent upon the shapes of the structuring elements selected and that the information extracted by an erosion sequence describes only a limited amount of information about the image. A useful analogy can be made between describing geometrical shapes by erosion sequences and the description of a random variable by its central moments. In general, to completely characterize a random variable, an infinite number of its moments are required. Likewise, an infinite number of erosion sequences, each specifying a unique set of structuring elements, will generally be needed to completely characterize a geometrical shape. Each unique erosion sequence would provide some shape-conditional information pertaining to the image eroded. In a limited number of situations (for example, where very little information is known about the classes of objects), the use of arbitrary and/or random (Ref. 24) erosion sequences may well be a viable technique to obtain features for unsupervised clustering tasks. However, when a priori information concerning the

geometrical characteristics of the targets, clutter and noise is available, there is little need for a "black art" erosion sequence selection process. Specifically, a large portion of Chapter IV will describe how to carefully select structuring element erosion sequences based upon a priori knowledge of the random and nonrandom information an image may contain. This, when combined with an awareness of the limits associated with geometrical measures such as erosion sequence image characterization, will provide several useful insights and techniques applicable to the algorithm selection process of cellular logic shape classification processors in general. The following section describes how combinations of erosions and dilations can be used for removing geometrical distortions and noise from binary images.

Two-Dimensional Closure and Opening

As described earlier, the dilation operation generally increases the size of objects (sets of state "1" cells) in binary images and, as a result, may not be as useful for image characterization as erosion sequences which erode (extract) information from an image. However, the erosion sequence will be shown to be quite sensitive to additive noise which may exist either as holes (state "0" cells interior to a group of state "1" cells) or concavities (one or more state "0" cells which protrude into the edge of a state "1" object) in the image. As a means of clarifying these noise descriptive terms (i.e. holes, concavities, and convexities), an analogy can be made to the geographical entities of a continent (an object of state "1" cells) which is surrounded by an ocean (field of state "0" cells). Continuing the analogy, a hole (a set of state "0" cells) in an object is analogous to a lake interior to the continent, a concavity would be analogous to an inlet or bay of state

"0" cells, and a convexity would be analogous to a peninsula of state "1" cells extending into the ocean of state "0" cells. The remainder of this section will describe the two-dimensional operations of closure and opening and to evaluate their potential for removing holes, concavities, and convexities from noise corrupted images.

Individually, erosion and dilation each exhibit some noise removal properties but in doing so, distort the image. The operations of opening and closure are an attempt to exploit the noise removal properties of the erosion and dilation operations. For example, Fig. 14(a) provides a noise free image labeled (H), two symmetrical structuring elements, and a noise corrupted version of the image (H) which is labeled (A) and exhibits noise in the form of holes, concavities, and convexities. Figure 14(b) illustrates the erosion of the noise corrupted image (A) of Fig. 14(a) using the structuring element B_1 . The shaded cells in these illustrations indicate the cells which were in state "1" prior to the indicated operation. Note that the eroded image $E_{B_1}(A)$ consists of only the three positions at which the structuring element B_1 could fit inside the image (A). Thus, one might conclude that the erosion operation removed the convexities and reduced the scale of the object. On the other hand, Fig. 14(c) illustrates that the dilation of the image (A) using the same structuring element filled in the hole and a concavity but in general left the convexities intact and increased the scale of the object. In an attempt to use the noise removal features of both dilation and erosion without changing the scale of the image, simple sequences of dilations and erosions have been applied (Ref. 11) to noisy images to remove (filter) the geometrical distortions caused by noise. In particular, an erosion followed by a




Fig. 14. Examples of Opening

dilation using the symmetrical structuring element B is defined as the opening of the image (A) with respect to the structuring element B,

$$A_B = D_B(E_B(A)). \quad (24)$$

Likewise, the closure of the image (A) with respect to the structuring element B is defined as

$$A^B = E_B(D_B(A)). \quad (25)$$

Each of these noise filters will be illustrated and discussed and observations will be made concerning their utility.

The opening of the image (A) using the structuring element B_1 of Fig. 14(a) is illustrated in Fig. 14(c) and labeled A_{B_1} . This result was obtained by dilating the state "1" cells of the image in Fig. 14(b). Note that the opening A_{B_1} is smaller than the original noise free image labeled H. Thus, due to the hole and concavities introduced by the noise, the opening of the noisy image resulted in an image which is significantly reduced in scale. To illustrate that the opening operation is sensitive to the geometry of the selected structuring element, observe that the erosion of (A) using the structuring element B'_2 is the null set because B'_2 could not fit into the noisy image. Thus, the dilation of the eroded image $E_{B'_2}^1(A)$ (i.e. the opening $A_{B'_2}$) is also the null set. The reader should observe that the erosion of the noise free image H using either structuring element (B_1 or B_2) results in a square four cell image and that the opening of H using the B_2 structuring element results in an image equal to H. An image which exhibits the property

$$A_B = A \quad (26)$$

is said to be smooth with respect to opening by the structuring element B .

The application of closure operations to the noise corrupted image (A) of Fig. 14(a) is illustrated in Fig. 15. The noise free image H , the symmetrical structuring elements B_1 and B_2 , and the noisy image (A) are repeated in Fig. 15(a), and Fig. 15(b) illustrates the closure of (A) using the structuring element B_1 . Note that the initial dilation operation has filled in the hole and the other missing state "1" cells but has increased the scale of the image. Therefore, the erosion of this dilated image (i.e. the closure A^{B_1}) is relatively similar to the original noise free image with the exception of the noise related convexities. For comparison purposes, the dilation and closure of the image (A) using the larger structuring element B_2 is illustrated in Fig. 15(c). Observe that in this example, the images A^{B_1} and A^{B_2} are identical except that the two noise related convexities on the left side of the image have been joined (the concavity formed between the two noise convexities has been filled in) by closure using the larger structuring element B_2 . In a similar manner, the closure of the noise free image H results in an image equal to H for either the B_1 or B_2 structuring elements. Therefore, the image H is smooth under closure with respect to either structuring element because

$$H^B = H . \quad (27)$$

Recalling how Eq. 14 provided an alternative method of evaluating the erosion operation of Eq. 13, there are alternate methods of evaluating the opening and closure definitions of Eqs. 24 and 25. For example, an




Fig. 15. Examples of Closure

equivalent description of opening is

$$A_B = \sum B'(a_{i,j}) \forall B'(a_{i,j}) \leq A. \quad (28)$$

Thus, the opening operation describes the set of cells spatially covered by the structuring element B' as it is translated vertically and horizontally interior to and contained in the state "1" cells of the image. Likewise, the closure operation can be described as the set of cells covered by the structuring element B as it is translated vertically and horizontally interior to and contained in the complement of the image. Thus,

$$A^B = \sum B(a_{i,j}) \forall B(a_{i,j}) \leq A^c. \quad (29)$$

Stated in terms of the uncomplemented image, closure is the set of cells not covered by the structuring element B as its root cell is positioned on every state "0" cell at which the structuring element does not intersect a state "1" cell of the image. The reader should compare these descriptions of opening and closure to the examples given in Figs. 14 and 15 and in particular observe that the only difference between A^{B_1} and A^{B_2} in Fig. 15 was the result of the ability of the von Neumann structuring element to fit into the concavity formed between the convexities on the left side of the noisy image (A) and the inability of the larger Moore structuring element (B_2) to perform the same task as the root cell of each structuring element was positioned on state "0" cells of the image.

Before directing the reader's attention to other classes of cellular logic operations, Fig. 16 provides a graphical description of the geometrical size relationships associated with the images formed by




Fig. 16. Closure and Opening Geometrical Relationships

the opening and closure operations. The set of events (binary images) in the domain is labeled \mathcal{Q} and is positioned between the opening operation on the left and the closure operation on the right. For simplicity of illustration, the larger indices imply larger binary images: the smallest image being the null image A_0 and the largest image being an all state "1" image labeled A_n . As the first step of the opening operation, the erosion of each event in the set of events \mathcal{Q} using the structuring element B' results in the set of eroded images labeled \mathcal{E} . Recalling earlier discussions, within each set of events \mathcal{Q}_n that erode to an event $E_n \in \mathcal{E}$, there exists an event $A_n^{\text{MIN}} \in \mathcal{Q}_n$ which is the smallest event that can erode to the event E_n . Exploiting the size ordering assumed for the events $A \in \mathcal{Q}$, the null set (A_0) is the smallest image that will erode to the null set E_0 (i.e. $A_0 = A_0^{\text{MIN}}$), the event A_4 is the smallest image that will erode to E_1 (i.e. $A_4 = A_1^{\text{MIN}}$), etc. The opening operation is then completed by dilating each event $E \in \mathcal{E}$ with the structuring element B . Referring to the definition of dilation (Eq. 4), it is apparent that the dilation of each event $E_n \in \mathcal{E}$ using the structuring element B will reconstruct the minimum event A_n^{MIN} . Therefore, the opening of an image,

$$A_B = D_B(E_B, (A)), \quad (30)$$

specifies (constructs) the smallest image that can erode to the erosion of the event (A). Thus,

$$A_B = A_n^{\text{MIN}} = \text{MIN}[A_j \in \mathcal{Q}_n] \quad (31)$$

where

$$\mathcal{Q}_n = \{A_j : E(A_j) = E_n, A_j \in \mathcal{Q}\}. \quad (32)$$

Using similar observations, the closure A^B of an image,

$$A^B = E_B(D_B(A)) , \quad (33)$$

constructs the largest event (A_{MAX}^m) which can dilate to the dilation of the image (A). Therefore,

$$A^B = A_{MAX}^m = \text{MAX} [A_j \in Q^m]$$

where

$$Q^m = \{A_j : D(A_j) = D_m, A_j \in Q\} . \quad (34)$$

While the relevance of these relationships may not be readily apparent, they do provide some useful bounds on the many to one mappings associated with opening and closure operations. In particular, it has been shown (Ref. 15) that

$$A_B \leq A \leq A^B . \quad (35)$$

In a more philosophical vein, frequent referral to Fig. 16 will be helpful in understanding the design and selection of two and three dimensional noise filters in the following chapter.

In summary, the closure and opening operations each have their unique capabilities and weaknesses, and the best operation to use for a given application will depend upon the type of noise one might expect to encounter (holes, concavities or convexities) and the operations which will follow the noise filter. For example, if some form of imagery was much less susceptible to holes and concavities than convexities, then the opening noise filter may be most appropriate. If, however, the data were equally susceptible to all three forms of noise, then the closure noise filter may be most useful since it generally removes two of the

three noise types. And, if immediately following the noise filter a series of erosions were used for shape or size discrimination (as is commonly implemented), then closure would generally be the most advantageous since, following the closure operation, any remaining convexities would be removed by the subsequent erosion sequence. Additional relationships concerning opening and closure operations are discussed in the next chapter but, as is always the case, the selection of an appropriate noise filter must be based on the physics of the problem (and solution) at hand. The following section describes special purpose two-dimensional cellular logic operations which will be used in the processor design of Chapter IV.

Two Special Purpose Transformations

The previous sections have described the two most common cellular logic operations (dilation and erosion) and two simple sequences of these operations (opening and closure). The definitions of these operations have been in terms of structuring elements as applied to binary images. This section will describe two additional transformations which apply multiple nonsymmetrical erosion sequences to images. While these transformations could be described in terms of complex sequences of structuring element operations, these operations are most descriptively defined in terms of neighborhoods and transition functions. Potential uses for the medial axis transformation (MAT) and the end erode operation (EERO) will be described, but specific applications will not be addressed until Chapter IV.

The Medial Axis Transformation. The medial axis transformation was first proposed by Blum (Ref. 1) as a feature extraction technique and many methods of its implementation are cited in Ref. 20. The medial

axis transformation will in general operate on an image to form a skeleton of that image: each point or cell of that skeleton being internal to the image and equally distant from the edges of the image. Thus the medial axis transformation is often referred to as a central axis transformation or a skeletonizing operation. Figure 17(a) defines the medial axis transformation used in this study, and Fig. 17(b) illustrates typical skeletons which are formed by the application of this MAT to three binary images. The state "0" cells represent the locations at which a state "1" cell has been changed to state "0" by the MAT and the remaining skeletons are indicated by a "1". Note that the skeletons are connected, are only one cell thick, and may contain one or more branches. Historically, skeletonizing operations have been used as a method of measuring an object's size or to normalize some uncontrolled variable of an image such as the variations of line widths in printed letters. More recently (Ref. 15) the MAT has been proposed as a data compression technique. In the following chapter, the MAT will be used in conjunction with an end erosion operation to extract additional shape information from images once a basic size discrimination process has been completed.

The End Erosion Operation (EERO). The end erosion operation is designed to detect cells of an image which are either isolated or connected to only one other cell and to remove (erode) these cells from the image. In general, the EERO is applied as a noise removal technique or as a method of detecting the ends of line segments (Ref. 3). The EERO operation is defined in Fig. 18(a) and its application is illustrated in Fig. 18(b) as applied three times to the skeletons of Fig. 17(b). In Fig. 18(b), the state "1" cells of the skeleton which

TRANSITION FUNCTION # 1

X	X	
X	Z	Y
	Y	

N_1

	Y	
Y	Z	X
	X	X

N_2

	Y	
X	Z	Y
X	X	

N_3

	X	X
Y	Z	X
	Y	

N_4

Change the state of cell Z from state "1" to state "0" if its neighbors labeled X are in state "0" and its neighbors labeled Y are in state "1". Otherwise, the state of cell Z remains unchanged.

TRANSITION FUNCTION # 2

	Y	
	Z	
X	X	X

N_5

X	X	X
	Z	
	Y	

N_6

X		
X	Z	Y
X		

N_7

		X
Y	Z	X
		X

N_8

Change the state of cell Z from state "1" to state "0" if its neighbors labeled X are in state "1" and its neighbor labeled Y is in state "0". Otherwise, the state of cell Z remains unchanged.

APPLICATION

- (i) Apply transition function # 1 to each cell of the matrix using neighborhoods N_1 , N_2 , N_3 and N_4 sequentially until the transitions within the matrix have stabilized.
- (ii) Apply transition function # 2 to each cell of the matrix using neighborhoods N_5 , N_6 , N_7 and N_8 sequentially until the transitions within the matrix have stabilized.
- (iii) repeat (i) and (ii) until both sequences are stable.

(a)

```

0 0 0 0 1 0 0 0 0      0 1 0      0 0 0 0 1
0 0 0 0 1 0 0 0 0      0 0 1 0 0      0 0 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1      1 1 1 1 1      0 0 1 1 0 0 0
0 0 0 0 1 0 0 0 0      0 0 1 0 0      0 0 1 1 1 0 0 0
0 0 0 0 1 0 0 0 0      0 1 0      1 1 1 0 0 0
                                0 1 0

```

(b)

Fig. 17. The Medial Axis Transformation (MAT)

C



A	B	C
D	E	F
G	H	I

C

NEIGHBORHOOD STATE CONFIGURATION	NEIGHBOR STATE *								
	A	B	C	D	E	F	G	H	I
1	0	0	0	0	1	0	0	0	0
2	0	0	#	0	1	#	0	0	#
3	0	0	0	0	1	0	#	#	#
4	#	0	0	#	1	0	#	0	0
5	#	#	#	0	1	0	0	0	0

* Neighbors which indicate state "#" can be in either state "1" or state "0". However, simple checks are made on these cells for each state configuration to insure that connectivity within sets of state "1" cells is retained.

(a)

[illegible]

(b)

Fig. 18. The End Brode Transformation (EABRO)

were eroded (nibbled away) by the EERO are indicated by the letter "E".

The end erosion operation can be used for several additional purposes such as counting the number of cells in a line segment, removing line segments from an image which are shorter than an arbitrary number of cells in length, or to effect the separation of long single branch line segments from line segments which exhibit multiple branches. It is this latter task for which the EERO will be exploited in the following chapter. However, before proceeding to Chapter IV, the concepts of two-dimensional cellular logic operations will be extended to three-dimensions.

Three-Dimensional Cellular Logic Operations.

The use of local spatial operations in three or more dimensions was documented by Unger (Ref. 25) in 1958. As described in Chapter II, the third dimension can represent intensity, range, height, or some other measure of information, and cellular logic operations can be readily applied to any of these types of data. Additionally, each two-dimensional concept or operation described in this chapter has an equivalent three-dimensional counterpart. For example, the concept of a structuring element extends to a three-dimensional set of neighboring cubes which describe a structuring volume with respect to a root cell. Likewise, while in two-dimensional binary space the state "1" cells are partitioned from the state "0" cells by their edge boundary, in multigrey level data, the surface described by the grey levels partitions three space into volume sets of cubic cells above and below the surface. Also, conveniently, the edge characteristics of concavities and convexities extend directly to surface characteristics of depressions (valleys) and protrusions (mountains). Holes, as used

with respect to two-dimensional data, do not exist in single value multigrey level data. If notation is established so that the cubes below the grey level surface are in state "1" and those above the surface are in state "0", binary volumes are established and a direct comparison can be made between binary two-dimensional cellular logic operations and their multigrey level counterparts. For example, two-dimensional binary erosion was earlier described as the root cell positions at which a two-dimensional structuring element could be contained in (i.e. fit into) an object/area of state "1" cells. In multigrey level data, a comparable binary erosion operation would specify the root cell positions (in three space) at which a three-dimensional geometric solid (i.e. a three-dimensional structuring element) could be contained in the state "1" cells of the binary volume below the surface. Thus, if a three-dimensional structuring element is defined as $B(a_{i,j,k})$ when positioned at cell $a_{i,j,k}$ of an indexed volume (A), then three-dimensional erosion is defined as

$$E_B(A) = \{a_{i,j,k} : B'(a_{i,j,k}) < A, a_{i,j,k} = 1, a_{i,j,k} \in A\} \quad (36)$$

where, as before, the symbol "<" requires that the three-dimensional structuring element B' be contained entirely within the volume of state "1" cells (i.e. does not protrude above the surface). Likewise, the dilation of the binary volume (A) with the structuring element B is the volume of cells spatially covered by the 3-D structuring element as its root cell is positioned at each state "1" cell in three-space. Therefore,

$$D_B(A) = \sum B(a_{i,j,k}); \forall a_{i,j,k} = 1, a_{i,j,k} \in A. \quad (37)$$

While in two-space, erosion generally removed convexities from edges and made concavities larger; in multigrey level data, the erosion operation will remove mountains from the surface and make valleys deeper and wider. Similarly, in binary two-space, dilation generally filled in edge concavities and exaggerated convexities, whereas in multigrey level data the dilation operation will fill in valleys and exaggerate mountains. Thus, conceptually, the three-dimensional operations of erosion and dilation perform noise filtering tasks quite similar to their binary two-dimensional counterparts.

As with two-dimensional operations, sequences of three-dimensional erosions or dilations generate larger equivalent three-dimensional structuring elements, and three-dimensional closure and opening operations perform similar noise removal tasks on three-dimensional surfaces. Figure 19 illustrates the application of a three-dimensional closure operation to a surface in three-space. Figure 19(a) illustrates a vertical slice (possibly a scanned line) of multigrey level data which contains both convexities and concavities. The closure operation using a cubic ($3 \times 3 \times 3$) structuring element is applied to the slice of data of Fig. 19(a) and the resulting slice is illustrated in Fig. 19(b). The dashed lines of Fig. 19(b) illustrate the results of the dilation portion of the three-dimensional closure operation. Note that the dilation operation has vertically biased the surface (slice) by the thickness of the 3-D structuring element (one cell), has removed the concavities which were smaller than the three cell width of the 3-D structuring element, and has thickened (exaggerated) both the large and small convexities. Following the subsequent erosion operation, the solid line of Fig. 19(b) illustrates the completed three-dimensional




Fig. 19. Three-Dimensional Closure and Opening

closure of the scan line in which the bias has been removed, the convexities have been returned to their original (pre dilation) size, and the large concavities (the concavities within which the 3-D structuring element could fit into) remain unchanged. Figure 19(c) illustrates the results of applying a three-dimensional opening operation to the scan line of data in Fig. 19(a). The dashed line indicates the results of the intermediate (erosion) step, the solid line indicates the completed opening, and the dotted lines the convexity removed by the opening operation. In addition to closure and opening, a third cellular logic operation is often considered for noise removal in multigrey level data. This filter is a local averaging function in which each cell's grey level is added to those of its eight nearest neighbors and the average value of this sum is the new value for the central cell. Figure 20(a) defines this local average noise filter, and Fig. 20(c) illustrates the results of applying this to the scan line slice of data illustrated in Fig. 20(b). The 3-D closure filter offers an advantage over the local average 3-D filter in that it has the potential to completely remove thin lines of correlated noise which is often experienced in real world line scanning sensors. Figure 21(a) illustrates a hypothetical example of a flat surface which, when scanned, exhibited a missing or badly biased line of data. Figure 21(b) illustrates the data of Fig. 21(a) following a local averaging noise filter. Note that the depression was smoothed somewhat but that the line of noise still remains obvious. Figure 21(c) provides the results of applying a 3-D closure operation to the data of Fig. 21(a) using a cubic (3x3x3) structuring element. While the 3-D closure operation totally removed the single line of noise, one must be cautious to select


The local average function, $\bar{A} = \text{AVG}(A)$, as applied to the i, j th cell of a two-dimensional array, $A = \{a_{i,j}\}$, is defined as

$$\bar{a}_{i,j} = \frac{1}{9} \sum_{k=1,3} \sum_{l=1,3} a_{i-k+2, j-l+2} ,$$

where " \sum " specifies arithmetic addition.


(a)

A slice of multigrey level data (A)




(b)

AVG(A)



(c)

Fig. 20. The Local Average Function



C

Fig. 21. 3-D Closure and Average Comparison

(a noise filter (and the corresponding structuring element) that does not remove or significantly distort small convexities or concavities which are necessary to the target identification task. In general, the structuring element (or equivalent structuring element) must be smaller than the smallest concavity or convexity one desires to retain as viable information in the data (surface or edge). This subject will be discussed at greater length with respect to the specific features and targets selected in the next chapter. As stated before, the physics of the random and nonrandom nature of the data and features must be understood in order to intelligently select a reasonable three-dimensional noise filter.

This completes the description of the two-dimensional and three-dimensional cellular logic operations which will be considered for use in the shape classification processor of Chapter IV. As described earlier, the emphasis of this chapter was to convey a geometrical understanding of several basic cellular logic operations and not to dwell on proofs or demonstrations which are generally straight forward and can be found in several sources (Ref. 1,8,11,15). Only "armed" with an understanding of the operations and their geometrical interpretations will the essence of the following become apparent.

C

IV. The Processor Design

This chapter provides a detailed description of how a cellular logic processor may be designed to perform an autonomous target classification task. Specifically, three hypothetical tactical targets are described, and these targets are assumed to be scanned on a flat surface by a nadir viewing range measuring sensor. Features which measure the true geometrical shape information contained in the three-dimensional (3-D) data are then selected, and noise and geometrical distortions associated with these features inspire the choice of noise filters. A unique design procedure is then developed that enables the designer to select, in a "cookbook" manner, the sequence of neighborhood transformations which will satisfy a specific set of performance goals. An overall design of the target classification algorithm is then developed and the false alarm rate associated with this processor design approach is discussed.

The Design Approach

A review of the literature (see Chapter II) clearly confirms that optimization techniques which rely on modeling input-output relationships and noise sources and applying analytical cost functionals have not been successfully applied to neighborhood transformation processors. While some progress has been made in constructing modeling tools for neighborhood transformation (Ref. 15) and statistically bounding simple geometrical properties of some specific transformations (Ref. 9), the nonlinear nature of these transformations and the difficulty in defining the signal (i.e. what is "patternness") in a form suitable for analytical modeling have limited the application of

classical optimization techniques to pattern classification tasks.

Recognizing these difficulties, but not being satisfied with a totally ad hoc design approach, an original systematic search technique of algorithm optimization is proposed. This technique negates the need to model the neighborhood transformations by applying a systematic search through all sequences of neighborhood transformations that could support a specific feature measurement requirement and evaluating each sequence's effectiveness via a unique performance oriented evaluation criteria. The characteristics of performance which form the basis for the proposed evaluation criteria are the probability of detection (P_d ; the probability of properly classifying a target), the probability of false alarm (P_{fa} ; the probability of improperly classifying a nontarget object as a target), the probability of misclassification (P_{mc} ; the probability of assigning a target to the wrong target class), and the probability of miss (P_m ; the probability of not detecting a given target). Thus, while the selection of features and feature extraction/measurement approach can be considered ad hoc (even though well justified in a scenario, resolution, and geometrical sense), the algorithm selection process is a rigorous sequence of step-by-step procedures which can be applied as an algorithm optimization technique to many feature - feature measurement pattern recognition tasks.

The design approach implemented can be summarized by the following sequence of design steps.

- a. Select two distinctive features which are common to each target class.
- b. Develop a generic approach to extract and measure these features that can be implemented by neighborhood transformations.

c. Using Monte Carlo simulations to accommodate geometric variables, systematically apply all neighborhood transformations (that could reasonably implement the feature measurement approach) to digitally synthesized noise free targets.

d. Develop a unique evaluation criteria which directly relates neighborhood transformation effectiveness to the performance goals of high P_d and low P_{fa} and P_m .

e. Determine which sequence of neighborhood transformations best satisfies the feature measurement task by applying the evaluation criteria to the results of the neighborhood transformation search.

f. Apply a likelihood ratio test to the feature-class conditional probability density functions (obtained from the Monte Carlo simulation of c. above) to partition the feature space into accept-reject regions for each feature and target class.

g. Define a strict class decision criteria which supports low P_{mc} and an alternate criteria which improves P_d by accommodating similarities of features between target classes. While not an integral part of the design procedure, three promising noise filters were evaluated for their deterministic distortive effects on geometrical objects and their ability to remove/smooth additive noise.

Finally, as in classical signal detection and estimation, the design goals of high P_d and low error rates are antagonistic. Thus, there were several instances during the design process that required a design decision that would support only one of the two goals. When such a conflict arose, the design approach most conducive to low false alarm rate was generally selected. This decision was based upon the fact that the lack of an accurate clutter (nontarget object) model prohibits a

realistic evaluation of the design's false alarm rate. Therefore, only by biasing the design toward a low false alarm rate could a bound on the processor's error performance be established. For design purposes, it was assumed that clutter was uniformly distributed over the features selected (i.e. over the feature space). In summary, this conservative design approach will provide a well understood baseline design against which future design decisions can be based. However, before describing the processor design, the scenario and other assumptions associated with the targets, environment, scanner, and noise will first be presented. One must be cautioned not to assume that the selection of appropriate noise filters, features, or feature extraction techniques are independent procedures as the chapter outline may hint. The various design decisions are closely related, not only to the extent that they influence each other, but also in that they are target, noise, and scenario dependent. Throughout the following sections this inter-relatedness will be described in detail; for herein lies much of the "learning curve" associated with cellular logic processor algorithm design.

The Scenario and Assumptions

This section will describe the scenario and assumptions associated with the targets, sensor, the environment, and noise.

The Targets. The targets are hypothetical and proportioned, as illustrated in Fig. 22, to represent three classes of tactical vehicles such as tanks or armored personnel carriers. The base of each vehicle model is composed of a rectangular shoe-box shaped object and the tops are modeled as horizontally truncated cones, the sides of which form 60 degree angles with the horizontal plane. The length, width, and height




Fig. 22. Target Class Models

of the bases and the height and base diameter of the tops are parameters which vary between target classes. The units of length associated with these target models can be described as "cell diameters" which are directly related to the altitude and instantaneous field of view (i.e. the spatial resolution) of the sensor.

The Sensor. The sensor is assumed to be an active range measuring device which unambiguously estimates vertical distance. Figure 23 illustrates the geometries of a line scanning sensor which is designed to collect a two-dimensional matrix (image) of range data. The sensor platform is translated in the y direction at a constant velocity normal to the direction of scan (the x direction) without over or under sampling. An ideal spot model is assumed where each range measurement provides the vertical (nadir) range as measured from the sensor to the center of its instantaneous field of view on the earth's surface. While the assumption that each range measurement provides information from a nadir viewpoint is not physically realizable, it is a very good approximation to range data which is collected within a few degrees of vertical where shadowing will not be prevalent in the range measurements. Thus, the multigrey level range measurements are made available to the target classification processor as the third dimension of a regular rectangular tessellation.

The Environment. The terrain, upon which the targets shall be randomly positioned and over which the scanner shall (conceptually) pass, is assumed flat. While this assumption (of flat terrain) is in general unrealistic, it was felt that presentation of the target models on randomly oriented surfaces would interfere with exploring more basic sensitivities to parameters such as noise, platform instabilities, and




Fig. 23. Conceptual Line Scanning Scenario

scale variations. Atmospheric anomalies such as fog, rain, snow, smoke, cloud attenuation, and scattering are not specifically modeled. However, since the noise model described in the following section was obtained by observing the randomness of real data, at least some uncalibrated amount of atmospheric noise/attenuation has been accommodated in the 3-D noise model.

The 3-D Noise Model. A brief analysis of two sources of three-dimensional infrared range data was used as a basis for the 3-D noise model to be described as normally distributed and white (independent) in time/space. The data samples (each sample consisted of about 400 data points) were all observed to be unimodally distributed, but the normality of these distributions (as measured by a Chi Square test) varied widely between the data samples. Additionally, the normality of the sample distributions varied inversely with the reflectivity of the various surfaces. These observations may be more an analysis limitation than a scientific observation because the higher reflective surfaces will generally provide improved range estimates and, since the truth model of the terrain was not exact, the samples which had nongaussian distributions could quite possibly be due to accurate measurements of nongaussian terrain surfaces. There was also a tendency for the data to be correlated for a distance of one or two cells in the direction of scan (the x direction) and uncorrelated between scan lines (the y direction). This was not unexpected since, in line scanning sensors, the time between line samples is typically much longer than the cell-to-cell sample times within a scan line of data. Thus, since the data exhibited an obvious unimodal tendency and was definitely normally distributed in several samples, a Gaussian noise model $N(0, \sigma^2)$ was

adopted. The white nature of the noise was initially chosen to be the worse case (uncorrelated vs correlated noise) and was later verified to be so during Monte Carlo simulations. The 3-D noise model does not incorporate clutter or other nontarget objects because insufficient data were available to aid in the design of such a model. However, an arbitrary measure of false alarm rate is applied to the processor as part of the Chapter V performance analysis.

In summary, the scenario consists of an airborne nadir viewing range measuring sensor which is propelled at a constant velocity over a flat terrain surface upon which randomly translated and rotated targets are positioned. The format of the collected 3-D range data is three-dimensional in that the third dimension grey levels provide a discrete topographical description of the flat terrain and target. The next section will now describe the architecture of the proposed processor.

The Processor Architecture

The cellular logic processor will consist of six parallel branches. Each branch is designed to measure incoming 3-D data for the existence of a specific geometrical feature associated with one of the three targets. Figure 24(a) illustrates the generic operations performed in a typical branch, and the design and purpose of each operation will be described in detail later in this chapter. Note that the operations of Fig. 24(a) have been grouped into three functional areas; noise filtering, feature extraction, and the decision criteria. Using these three functional areas, Fig. 24(b) illustrates how six branches are combined to form the processor. While this description implies that all the 3-D data is operated on by all six branches and that the operations




Fig. 24. The Processor Architecture

within each branch are serially implemented, the actual physical mechanization of a cellular logic processor may not necessarily follow this approach. However, the parallel-serial architecture of Fig. 24(b) does accurately portray the functional design and, for clarity of presentation, was the approach selected for implementation on a general purpose digital computer (CDC 6600).

Feature Selection

Two features will be described for each target class. Since the information content of the described data is unique in that it is truly three-dimensional, the features are specifically selected to measure three-dimensional target shape information.

Figure 25(a) illustrates a sample of a class 3 target scanned and digitized in 3-D range data format. Rather than selecting gradient or texture features as is commonly done with passive data, a more ancient and simple technique of thresholding (not currently in favor with the image processing community because of the diurnal variations evident in passively collected imagery) will be used. The 3-D image of Fig. 25(a), for example, would be thresholded at two levels; once at one-half the height of the base, and a second time at one-half the height of the top above the base. The thresholding operation is defined so that cells in the threshold plane are assigned state "1" if the center of the cell lies above the plane and state "0" otherwise. The results of thresholding the top and base of a randomly positioned class 3 target is illustrated in Fig. 25(b). The geometrical distortions exhibited by the digitized and thresholded binary images are commonly referred to as digitization noise and, because digitization noise significantly influences the thresholded binary images, any proposed shape




Fig. 25. A Digitized and Thresholded Class 3 Target

classification algorithm must be relatively insensitive to this noise source. Using the thresholds as defined above, the two resulting features (2-D binary images) provide a three-dimensional measure of the targets or any other objects on the terrain. This measure, of course, does not totally characterize the geometrical shape of an object, and very little imagination is required to describe nontarget objects which, when thresholded, provide similar 2-D binary images. Ideally, however, very few nontarget objects would match both binary images at the thresholds selected for the class 3 targets. While it is possible to more completely characterize three-dimensional objects by using additional height thresholds, additional features such as corners, edges, and curves, or using three-dimensional geometrical volume measurements (similar to the 2-D binary image area measurements described in the previous chapter), the decision was made to fully exploit the unique qualities of actively collected three-dimensional range data by limiting the number of features to two and keeping the features extremely simple.

Additional forms of geometrical distortions and noise which are typical to this form of imagery data will next be described and appropriate noise filters proposed and discussed. One potential distortion is introduced by the noise filters themselves (much as a low pass filter smooths off rising and falling edges). Thus, the 3-D and 2-D noise filters must be specified and their deterministic effects on the selected features evaluated prior to feature extraction/measurement algorithm selection.

The Noise Filter

This section will discuss the selection of the 3-D closure

operation (vs local average or 3-D opening operations) as the noise filter and describe how the choice of an appropriate 3-D structuring element for use in the 3-D dilation-erosion (closure) sequence is related to the size and shape of the smallest topographical feature of interest. It will also be shown that the selected 3-D noise filter directly impacts the need for and selection of an appropriate 2-D noise filter.

The 3-D Noise Filter. Under present assumptions, noise free three-dimensional data representative of a scanned target would consist of a multigrey level surface which contains two convexities; one which is large and formed by the digitization of the randomly oriented shoe-box base and a much smaller convexity formed by the digitization of the truncated cone which is centrally superimposed on top of the base. In order to evaluate how a 3-D closure or opening operation would deterministically influence a noise free topographical surface such as this, one should recall (from the previous chapter) that a 3-D closure noise filter would, in general, remove small concavities (i.e. fill in potholes within which the 3-D structuring element could not fit down into) but leave convexities on the surface unchanged. Conversely, the 3-D opening operation would remove small convexities (i.e. clip off bumps on the surface within which the 3-D structuring element could not fit up into) while leaving concavities in the surface intact. Thus, the use of a 3-D opening operation as a noise filter limits the size of allowable 3-D structuring elements (and, hence, limits the size of convex noise spikes that can be filtered) to those which can fit up into the smallest convex 3-D feature (the top of target class three) or else part or all of the convex feature would be removed by the noise filter.

Since the selection of a 3-D structuring element for use in a 3-D closure operation is not limited in this manner, the 3-D closure noise filter would appear to be less limited in design flexibility and hence more desirable. However, the earlier assumption that noise is statistically independent implies that a large 3-D structuring element is not required to remove noise since noise will "most often" appear as independent spikes on or small potholes in the 3-D surface. Additionally, while a 3-D closure operation which used a large 3-D structuring element may work quite well under the clutter free environment assumed for this analysis, in a less academic environment a 3-D closure operation which uses a large 3-D structuring element would tend to connect the target to other nearby surface convexities (trees, buildings, other vehicles or terrain) if this large 3-D structuring element could not fit down into the concavity which separates the target from other convex objects. The following example will attempt to illustrate these concepts by describing the 3-D closure and opening operations as applied to a convex object in the presence of noise.

Figure 26(a) illustrates a 3-D surface realized by digitizing the height of a shoebox shaped object, the sides of which are aligned with the axis of a regular rectangular scanning grid. When using the 3-D structuring element B_1 (of Fig. 26(a)), neither a 3-D closure nor a 3-D opening operation will change (distort) the noise free surface labeled H. This is because structuring element B_1 can fit down onto/around (closure) and up into (opening) each cell above and below the surface respectively. Figure 26(b), on the other hand, illustrates the results of applying a 3-D closure and opening operation to this same surface using the seven cell 3-D structuring element labeled B_2 . Note that the

AD-A111 326

AIR FORCE INST OF TECH WRIGHT-PATTERSON AFB OH F/G 17/9
THE DESIGN AND PERFORMANCE CHARACTERISTICS OF A CELLULAR LOGIC --ETC(U)
APR 81 L A ANKENY
AFIT/DS/EE/81-1

UNCLASSIFIED


NL

2.2

1.104



END
DATE
FILMED
4 82
DTIC






Fig. 26. Closure and Opening Applied to A 3-D Image

concave corners of the surface have been filled in by the closure operation because the 3-D structuring element B_2 could not fit down into these cells without intersecting the surface. Likewise, during a 3-D opening operation, the convex corners of the convex object were removed because the 3-D structuring element B_2 could not fit into these cells. Since the 3-D opening operation reduced the volume of the convexity significantly, when thresholded at one half the height of the original shoe-box object (1.5 cells), its 2-D binary image, $Th(H_{B_2})$, is smaller than $Th(H^{B_2})$ or $Th(H)$.

Continuing the example, Fig. 27(a) illustrates the 3-D surface (H) of Fig. 26(a) following its distortion by three kernels of additive noise. A central vertical section of this noise corrupted surface is also illustrated in Fig. 27(a) and the concavity and convexities formed by the additive noise are labeled n_1 , n_2 , and n_3 . Observe that the concavity n_2 and the convexity n_3 have created a hole and an unconnected state "1" cell respectively in the thresholded image $Th(A)$. If the surface labeled (A) of Fig. 27(a) is filtered by a 3-D closure operation which uses the 3-D structuring element B_1 (of Fig. 26(a)), the surface and thresholded image of Fig. 27(b) are obtained. Note that the concavity n_2 has been filled in (because the structuring element could not fit down into it) and that while the convexity n_1 remains unchanged, the convexity n_3 has become connected to the shoe-box object because the structuring element could not fit into the concavity formed between the convex noise spike n_3 and the larger shoebox shaped convexity. The application of a 3-D opening operation to the noise corrupted surface of Fig. 27(a) results in the surface labeled A_{B_1} illustrated in Fig. 27(c). Note that the 3-D opening operation has not




Fig. 27. Closure and Opening Applied to A Noise Corrupted 3-D Image

only removed the convex noise spikes (n_1 and n_3) but has also removed a large portion of the larger convexity because the concavity n_2 has prevented the 3-D structuring element (B_1) from fitting all the way up into the shoe-box shaped convexity. The resulting thresholded image $Th(A_{B_1})$ is the null set.

To conclude this example, Figs. 28(a) and (b) illustrate the result of applying 3-D closure and opening operations to the noise corrupted image of Fig. 27(a) using the seven cell 3-D structuring element labeled B_2 in Fig. 26(a). Note that the results of using this structuring element are very similar to those obtained using the 27 cell structuring element B_1 . In each case, however, the thresholded binary images more closely match the ideal (noise free) threshold image of Fig. 26(a) when the 3-D closure and opening operations use the seven cell B_2 structuring element. This of course will not always be the case. However, under the present assumptions and for small convex features of interest, the 3-D closure operation using the B_2 structuring element will generally be more useful than a 3-D opening operation because it will retain sufficient convexity to provide a 2-D threshold image while a single noise concavity could cause a 3-D opening noise filter to erode the available information to heights below the threshold value. Therefore, even though the 3-D closure noise filter can potentially enlarge portions of convex features and connect nearby objects to targets of interest, the 3-D closure operation retains all of the available convexity information of the noisy 3-D data as well as fill in small concavities which may appear in the surface. Thus, following a 3-D closure operation, the thresholded 2-D image will at least contain the thresholded 2-D image obtained by thresholding the unfiltered 3-D




Fig. 28. Closure and Opening Using Structuring Element B_2

surface. This relationship can be described as

$$\text{Th}(A) \leq \text{Th}(A^B) \quad (38)$$

where $\text{Th}(\)$ is the height threshold operation, A is a 3-D surface of multigrey level data, and A^B is the closure of the surface A using the 3-D structuring element B . Equation 38 is a direct extension of the two-dimensional relationship of Eq. 35 for a 2-D binary image consisting of state "1" cells in a state "0" field. In three space, the convention of state "1" cells being below the surface and state "0" cells being above the surface extends to the concept that the surface resulting from a 3-D closure operation will cover the unfiltered surface (i.e. the unfiltered surface will be below or equal to the 3-D closed surface). The results of thresholding surfaces sharing this relationship should be obvious.

A third potential 3-D noise filter earlier described as a local averaging function was not selected for use because, unlike 3-D closure, it does not exhibit the ability to totally remove noise which is either uncorrelated or correlated in one direction (as illustrated in Fig. 21) and a relationship similar to Eq. 38 cannot be shown for an unfiltered surface and its locally averaged counterpart.

In summary, the 3-D closure noise filter was selected because, in a noisy environment, it removed concavities in the 3-D surface without reducing the size of any convexities that may exist on the 3-D surface. This property of 3-D closure operations is most important since the target models appear as convex objects in three-dimensional height data. The 3-D closure operation will use the seven cell 3-D structuring element labeled B_2 in Fig. 26(a) because it is large enough to remove

independent or slightly correlated concavities in the surface as well as fill in occasional missing lines of scanned data. This same structuring element should be small enough to prevent unnecessary surface growth due to connecting the convex target objects to nearby convex noise spikes or other convex objects. The local average and 3-D opening filters were not selected because they did not exhibit the ability to totally remove concavities nor did they retain all the convexity information available within the unfiltered 3-D surface. The selection of this 3-D noise filter has been shown to be closely related to the noise characteristics, the size and nature of the smallest 3-D features of interest, and the method by which the 3-D feature is measured (i.e. thresholding). The following section will now describe how the 3-D closure noise filter, when combined with the erosion method of measuring thresholded 2-D images, directly influences the need for a 2-D noise filter.

The 2-D Noise Filter. This section will discuss how the geometries of the targets, the use of a 3-D closure noise filter, and a yet to be described method of feature measurement combine to negate the need for a 2-D noise filter operation.

In Chapter III, it was suggested that a 2-D closure operation could be used to fill in holes or edge concavities of state "0" cells in 2-D binary state "1" objects. For a 2-D binary object (obtained, perhaps, by thresholding the shoe-box shaped base of a target model) to contain a hole or edge concavity, one or more of the cells describing the top surface must be corrupted with noise to a height below the threshold value. If the noise free height of the i,j th cell of an object is $a_{i,j}$ and the threshold value selected is Th , then a zero mean additive noise

source (\tilde{x}) would induce a hole in the surface's thresholded image if the realization of noise $x_{i,j}$ had a value such that

$$c_{i,j} + x_{i,j} < Th. \quad (39)$$

The probability of a hole or concavity occurring at the i,j th position of a binary image formed by thresholding a noise corrupted 3-D surface is

$$P [c_{i,j} + x_{i,j} < Th] = \int_{x=-\infty}^{x=Th} f_{\tilde{x}}(x) dx = p \quad (40)$$

where $f_{\tilde{x}}(x)$ is the probability density function of the noise source. Recalling that a 3-D closure operation can be used to fill in surface concavities within which the 3-D structuring element cannot fit down into, the application of a 3-D closure operation to a 3-D surface insures that a hole or concavity will not exist in the thresholded 2-D image of that surface unless the 3-D structuring element fits down into the concavity on the 3-D surface to at least a depth below the threshold height. Therefore, if prior to thresholding, the 3-D surface were filtered using a 3-D closure operation that used a 27 cell ($3 \times 3 \times 3$) 3-D structuring element, the probability of a hole existing in the thresholded 2-D binary image would be p^9 assuming spatially independent noise. Likewise, the probability that any cell on the edge of the same 2-D binary image would exhibit a concavity would again be p for an unfiltered surface and no more than p^4 for a 3-D closure filtered surface. The probability of such an event occurring is, of course, finite and is related to the shape and size of the 3-D convexity within which the concavity may occur as well as the shape of the structuring

element. However, the application of a 3-D closure noise filter operation to the 3-D data prior to thresholding will generally make the probability of a hole or concavity occurring in the 2-D binary image very small, hence, negating the need for a 2-D closure noise filter.

The other potential 2-D noise filter discussed in Chapter III was the 2-D opening filter which was effective in removing edge convexities and small unconnected cells of state "1" noise. Since the 3-D closure operation was shown not to be effective in removing 3-D surface convexities which, when thresholded, form these convex forms of 2-D noises, the need for a 2-D opening noise filter appears quite logical and appropriate. However, the next section will shortly describe that the initial step used to measure the size and shape of the 2-D binary images is a sequence of one or more erosion operations. Thus, assuming the same 2-D structuring elements would be used, any convexities present in the 2-D images which would have been removed by a 2-D opening operation would also be removed by the shape measurement erosion sequence and a 2-D opening operation would therefore be redundant. In the event that 3-D features are selected (or noise is experienced) which invalidates any of these assumptions, then the geometries associated with this modified environment should be reviewed to determine 2-D noise filter requirements.

Feature Measurement

The circular and rectangular images obtained by height thresholding the scanned target models were selected as features because they provide simple but unique 2-D measurements of each target's 3-D geometrical volume. This section will describe an approach to measure these features. The circular feature will be measured by simple erosion

sequences while the more complex rectangular features require additional processing to extract shape information. Significant issues which directly influence the performance of the processor such as target translation and rotation, algorithm selection, processor architecture, and the probability of false alarm rate will be specifically addressed.

Circle Measurement. When thresholded, the truncated conical tops of each target class will result in circular 2-D disk shaped objects. However, due to the finite resolution of the digitizing grid (i.e. the scanning sector) and the assumed random positioning of the target on this grid, many 2-D binary image realizations can be experienced for each digitized target top. Several typical examples of thresholded target tops for each target class are illustrated in Fig. 29. This section will describe how a simple erosion sequence can be selected to provide a useful geometrical measure of digitized randomly positioned circular disks of arbitrary radius.

In Chapter III, each of the cells remaining in state "1" following an erosion sequence were shown to be the root cell positions at which the erosion sequence's equivalent structuring element fit into the state "1" cells of the uneroded binary image. Thus, by counting the number of residue cells (the state "1" cells which remain following an erosion sequence), a limited measure of size information relative to the uneroded image and the equivalent structuring element can be inferred. Figure 30(a) provides 2-D images obtained by digitizing three different objects. The residues obtained by eroding these objects with two different erosion sequences will illustrate the proposed size measurement technique as well as introduce the issues associated with erosion sequence selection. The erosion sequence B_1 , when applied to




Fig. 29. Typical Thresholded Target Tops




Fig. 30. Residue Count Relationships

each image of Fig. 30(a), results in the images (residues) illustrated in Fig. 30(b). The residue counts for these images are one, seven, and eleven for images A_1 , A_2 , and A_3 respectively. Stated equivalently, the structuring element B_1 (a 3×3 square) fit into the state "1" cells of image A_1 once, A_2 seven times, and A_3 eleven times. Note that the erosion residue of image A_1 implies that the largest area of the image A_1 is about the same size and shape as the structuring element B_1 . The larger residue counts of images A_2 and A_3 imply that the structuring element B_1 fit into these images in several positions but, other than knowing that the images A_2 and A_3 are larger than B_1 , very little can be concluded about their gross shape. Continuing the example, eroding the binary images of Fig. 30(a) with the structuring element B_2 results in the residues illustrated in Fig. 30(c). Note that the image A_1 could not contain the structuring element B_1 and therefore eroded to the null set, image A_2 eroded to a residue of one, and the image A_3 , which had a larger residue count than A_2 for the B_1 erosion sequence, eroded to the null set. Thus, a large count of residue cells does not always indicate a large object and, in general, larger structuring elements which fit into an object at only a few positions convey more about the gross uneroded shape of the object than a smaller structuring element could. Qualitatively, given the task of attempting to detect (measure) circular shaped binary objects of a certain diameter, one should select a structuring element (erosion sequence) which has a maximum width slightly smaller than the diameter of the circular object and, following similar logic, has a circular shape. Even with this knowledge, the large variety of images obtained by digitizing randomly translated disks makes the selection of an appropriate erosion sequence for each

truncated target top a nontrivial task. Stochastic Geometry (Ref. 9) has demonstrated some success in establishing statistical measures of geometrical questions such as this, but these measures are not directly related to specific geometric shapes of the digitized images. Therefore, a Monte Carlo simulation was designed to address the question of what structuring elements were best suited to fit consistently but not too loosely (i.e. which structuring element fit snugly) into digitized randomly positioned disks of various diameters in the absence of noise.

The Monte Carlo analysis was implemented by randomly positioning (uniformly) thirty-five circular disks of a given radius onto a digitizing grid. Since geometrical probabilities do not address specific geometric shapes, a sequence of trial and error experiments were used to establish that, for arbitrary random number (translation) generator seeds, the relative frequency of occurrence of specific geometric shapes (resulting from digitizing the randomly positioned disks) was repeatable if at least 35 disks were used for a given Monte Carlo experiment. The digitized 2-D binary images of these disks were then individually eroded by each of the first 12 structuring elements illustrated in Fig. 31 and the number of cells which remained following each erosion sequence (i.e. the residue count) was recorded. The radius of the disk was then increased by .1 unit and the procedure repeated for disks of radii between 1.5 and 5.5 units. The structuring elements of Fig. 31 were selected because they represent the smallest symmetrical erosion sequences which can be implemented by combining the basic von Neuman (SE1) and Moore (SE2) neighborhoods.

Recalling that the intent of the Monte Carlo analysis was to




Fig. 31. Symmetrical Structuring Elements

determine which structuring elements fit snugly into the digitized disks, the concept of fitting snugly had to be bounded. An obvious lower bound for the residue count of a snug fitting structuring element was established by requiring the structuring element to fit into each of the thirty-five events (randomly positioned digitized disks of a specific radius) at least once. An upper bound on the residue count for a snug fit must also be established to limit the number of positions (i.e. limit how loosely) a structuring element can fit into the digitized disk. The upper bound, however, cannot be arbitrarily specified since there are only a limited number of structuring elements available, and one structuring element cannot be declared too loosely fitting until the next larger structuring element starts fitting snugly. The Monte Carlo evaluation, therefore, provided a means to establish the size of the smallest circular disk within which each structuring element of Fig. 31 could consistently be contained at least once for all 35 events. Circular disk shaped objects of radius between 1.5 and 5.5 pixels were thereby partitioned into sets, each set specifying the size of disks, within which, one of the structuring elements of Fig. 31 was determined to be snug fitting according to the residue count criteria described above.

To illustrate the Monte Carlo results, Fig. 32(a) plots the smallest residue count experienced for the SE1, SE2, and SE3 structuring elements as applied to the 35 samples of digitized circular disks. As the radii of the circular objects were increased from 1.5 to 3.5 units, the point at which each structuring element begins to fit into the circular objects clearly specifies the lower and upper radii bounds within which each structuring element fits snugly. For example, SE1




Fig. 32. Circle Measurement Algorithm Selection

did not fit at least once into all 35 randomly translated and digitized disks until the disks attained a radius of 1.6 units. SE1 then continued to fit snugly until SE2 started fitting into all 35 disks of 2.1 units radius. Thus, if a circular object had a radius of 2.2 units, then the SE2 erosion sequence would fit snugly into this object while SE1 would fit too loosely and SE3 would be too large. Figure 32(b) summarizes these results by listing the radii over which the smaller 12 structuring elements of Fig. 31 were determined to be snug fitting. While these results establish which erosion sequence should be used to measure circular objects of a given radius based upon the snug fit criteria, it would also be desirable to understand how well these selected erosion sequences perform their shape measurement tasks. For example, when geometrically measuring discrete shapes by counting the number of cells which remain following an erosion sequence, it would be desirable to know that the measurement technique consistently provided the same residue count regardless of where the circular object was placed on the digitizing grid. As an aid in describing this concept, for a circular object of given size and a specific erosion sequence, the residue set is defined as the set of unique residue counts (integers) obtained when the continuous object is digitized and eroded at many random positions on a digitizing grid. Likewise, the residue set spread is defined as the number of integers separating the largest and smallest events in the residue set. Small nonzero residue sets and narrow residue set spreads would, therefore, be desirable since this would imply that the structuring element not only fit snugly but also fit consistently within the digitized images. Figure 32(c) illustrates the residue set spread function exhibited by structuring elements SE1, SE2,

and SE3 as they were applied to 35 samples of digitized circles of radii 1.5 to 3.5 units. More generally, the largest residue set spread exhibited by the snug fitting structuring elements as defined in Fig. 32(b) was four and was more typically only two or three integers wide. The desirability of these relatively consistent and narrow residue set spreads will be further discussed following the next section which describes the measurement of rectangular shaped binary objects.

Rectangle Measurement. This section will describe how an erosion sequence is selected and used in conjunction with a medial axis transformation (MAT) and an end erosion operation (EERO) to measure the size and shape of 2-D binary rectangular objects. The proposed design addresses the rotational sensitivity of discrete erosion sequences, the application of MAT and EERO to extract length and shape information, and how the goal of low false alarm rate is supported by the algorithm selection process.

The 2-D binary rectangular features obtained by thresholding the target bases are significantly more difficult to measure than circular features because they are not rotationally invariant and, since they cannot be geometrically described by a single measurement (such as radius), both length and width measures must be developed. As with circular features, the basic approach of selecting a structuring element which fits snugly into the rectangles (i.e. selecting an erosion sequence that consistently reduces the rectangles to a small nonzero number of state "1" cells) can be used to measure the rectangle's minimum dimension, the width. However, the erosion sequence selection process is complicated by the fact that the structuring elements are not round and their ability to fit snugly into a digitized rectangle is

highly dependent upon the uncontrolled orientation of the target (rectangle).

Figure 33 conceptually illustrates the relative sensitivity of two erosion sequences to the orientation angle of a rectangular object. If we assume that the first erosion sequence exhibits a round equivalent structuring element of radius $w/2$, the circular structuring element (illustrated in Fig. 33(a)) fits into the rectangle at any orientation and the locus of its center, which represents the residue of the erosion sequence in this example, provides a rotationally invariant measure of the rectangle's length. On the other hand, Fig. 33(b) illustrates that a second sequence of erosions which exhibits a square equivalent structuring element of width w can only fit into the same rectangular image at one orientation and, as a result, provides a rotationally sensitive measure of the rectangle's width and length. Unfortunately, given a specific rectangular object, the finite number of discrete symmetrical structuring elements (erosion sequences) to choose from does not generally allow the selection of one which is circular in shape and fits snugly into the digitized rectangular object at all rotations and translations. Figure 34 illustrates a typical example of how the limited number of symmetrical erosion sequences forces the system designer to use a less than ideal erosion sequence and how an additional operation can be implemented to overcome some of its inherent limitations. Figure 34(a) illustrates the results of digitizing a 6.4 by 12.0 pixel rectangular object at three different orientations with respect to the digitization grid. Of course, the original binary images consisted of all state "1" cells but, since the illustrated images have been eroded by the SE7 structuring element of Fig. 31, the eroded cells




Fig. 33. Conceptual Example of Erosion Sequence Sensitivity to Target Rotation

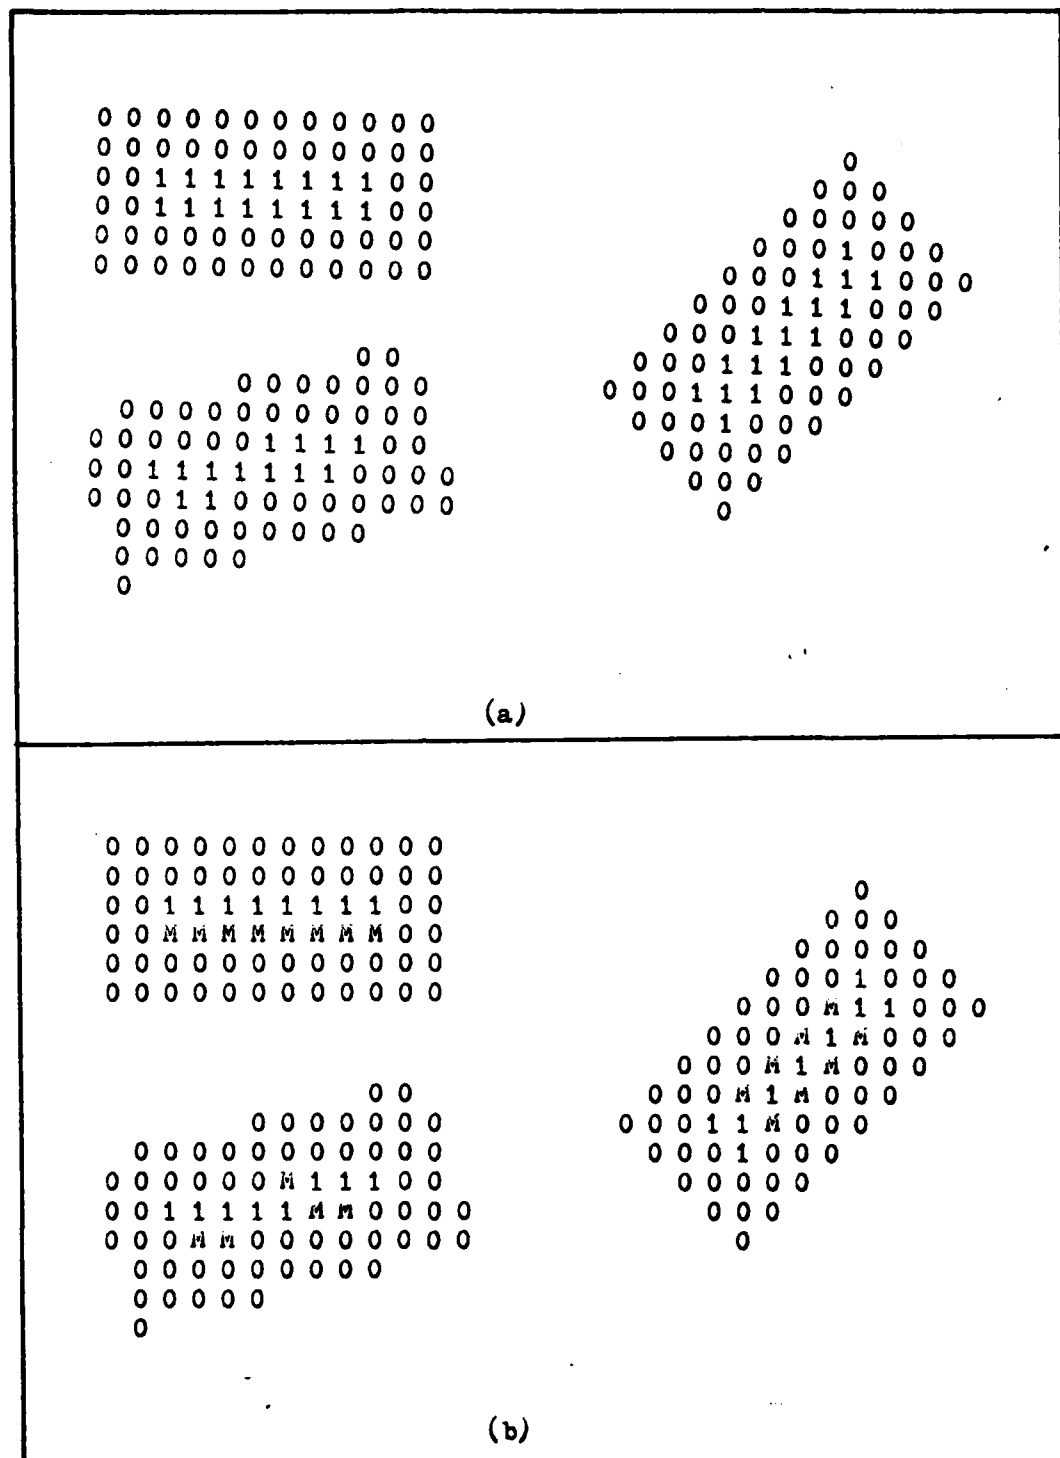


Fig. 34. Application of the Medial Axis Transformation

have been changed to state "0" and the remaining state "1" cells are the residue of the erosion sequence. Without illustration, it should be obvious that a smaller structuring element such as SE6 would fit into the original binary images at more positions (i.e. more loosely) than SE7 and is therefore not as snug a fit and is less desirable for rectangle width measurement. Likewise, the next larger structuring element (SE8) is too large to fit into the original digitized rectangular images at all translated and rotated positions (even though it does fit into the examples of Fig. 34) and is therefore not a viable snug fitting structuring element for this sized rectangle. Thus, for this sized rectangle, the SE7 erosion sequence must be used even though its residue (state "1" cells) is more than one cell thick and a simple cell count procedure cannot be used to estimate the length of the rectangle. While there is little that can be done about SE7's loose fit, the medial axis transformation (MAT), described in Chapter III, can be applied in situations such as this to reduce the residue cells to a skeleton one cell in width. Figure 34(b) illustrates the application of the MAT to the residues of Fig. 34(a). Note that the state "1" cells which were removed by the MAT have been changed to state "M" (for illustration purposes) and a simple cell count of the remaining state "1" cells can be used to estimate the length of the rectangle. Thus, the MAT will be applied in conjunction with the basic erosion sequence in each branch of the processor which is designed to measure the length and width of rectangular objects.

A Monte Carlo analysis was used to determine which structuring element, when used in conjunction with the MAT, demonstrated both a snug fit and an insensitivity to rectangular image translation and rotation.

In a manner similar to that used for the circular features, continuous rectangles of varying minimum dimension (the length was arbitrarily set to twice the width to prevent telephone-pole shaped rectangles which are not typical of tactical vehicles) were positioned on a digitizing grid at thirty-five random translations at seven equally spaced angles over the interval (0, 45) degrees. Each event (each digitized randomly positioned rectangle) was then eroded by each of the structuring elements of Fig. 31, skeletonized by the MAT, and the number of remaining cells was recorded. Upon reviewing the results of the Monte Carlo simulation, it became apparent that while a simple "snug fit" erosion sequence selection criteria (similar to that used for circular features) did specify which structuring elements were snug fitting with respect to the width measurement, this criteria did not establish which erosion-MAT sequences provided consistent rectangle length estimates. Fortunately, the residue set spread function (which was briefly described in the previous section) provides a great deal of information concerning the rotational invariance and length measurement consistency of an erosion-MAT sequence. Figure 35 provides the residue set spread functions obtained by Monte Carlo analysis for the typical SE6 and SE7 erosion-MAT sequences. For rectangle widths below 4.6 pixels, the residue set spread is zero for both sequences since neither structuring element fit into any of these events. As the width of the rectangles was incrementally increased, the structuring elements began to fit into the digitized images and increase in size since some events (digitized randomly positioned rectangles) can contain the structuring element at several positions but other events cannot contain it at all due to the rectangle's angle of rotation with respect to the digitization grid.

The duration and height of this initial relative maximum is directly related to a structuring element's shape. For example, a square structuring element such as SE6 will exhibit a wide and high relative maximum since its size (SE6's diameter) varies significantly as a function of rotation angle. On the other hand, a less square structuring element such as SE7 fits into all events shortly after it starts fitting into some and, therefore, its relative maximum is not nearly as large as the square structuring element SE6. A truly round structuring element would, of course, not exhibit any initial relative maximum at all. Continuing the example, both structuring elements start fitting at least once into all the events (an occurrence not commonly experienced) at 6.8 pixels, and it is at this point (circled in Fig. 35) that the residue set spreads begin to decrease. The sharp decrease in residue set spread exhibited by SE6 is due to the tendency of square structuring elements to fit into a rectangular object at many positions or not at all. Nonsquare structuring elements, on the other hand, exhibit a more gradual decrease in residue set spread and will, in general, exhibit a smaller and more lengthy initial relative minimum. Finally, as the rectangle width is increased, both structuring elements experience increases in residue set spread due to the increased number of residue cells which results in increased skeleton variations. Therefore, the residue set spread function provides both a measure of a structuring element's roundness as well as a consistency measure for the length estimate characteristics of the erosion-MAT sequence. In the example of Fig. 35, assuming for illustration purposes that no other erosion-MAT sequences are available, both SE6 and SE7 become snug fitting with respect to the rectangle's width at 6.8 pixels but, since




Fig. 35. Typical Residue Set Spread Functions

SE6 exhibits a more consistent estimate of rectangle length (i.e. has a lower residue set spread) over the interval of rectangle widths from 6.8 to 7.2 pixels, it would be the preferred erosion-MAT sequence in this interval. SE7 would then be the preferred erosion-MAT sequence for measuring rectangles with width greater than 7.2 pixels until a larger snug fitting structuring element exhibited an erosion-MAT residue set spread smaller than SE7. The preferred erosion-MAT sequences were established by Monte Carlo simulation for rectangular objects between 2.8 and 12.0 pixels in width by applying this same selection criteria to the set of structuring elements of Fig. 31. The results are provided in Table 1. Several erosion-MAT sequences do not appear in the tabulation because the square shapes of their equivalent structuring elements prevented them from attaining a narrower residue set spread than other snug fitting erosion-MAT sequences. In conclusion, for an erosion-MAT sequence to be preferred with respect to the measurement of rectangles of a given width, not only must the residue counts obtained during the Monte Carlo simulation be nonzero for all 35 events (i.e. snug fitting), the sequence must also exhibit the smallest residue set spread.

While the above proposed erosion-MAT sequence followed by a simple cell count provides a reasonable amount of shape discrimination capability, the many-to-one property of erosion sequences is a limiting factor because there are many nonrectangular objects which will exhibit skeletons with a cell count identical to that of any given rectangle. Figure 36(a) illustrates three geometrical shapes which have been eroded using the SE2 erosion sequence and skeletonized by the MAT. As in the previous example, the state "0" cells specify those state "1" cells of the original image which were eroded by the SE2 erosion sequence and the

TABLE I

PREFERRED EROSION-MAT SEQUENCES
FOR RECTANGULAR OBJECTS
(REF. FIG. 37)

STRUCTURING ELEMENT	RECTANGLE WIDTH *	
	FROM	TO
SE 1	2.8	4.0
SE 2	4.0	4.8
SE 3	4.8	5.3
SE 4	5.3	7.0
SE 6	7.0	7.2
SE 7	7.2	9.0
SE 11	9.0	10.5
SE 12	10.5	12.0

* PIXELS


```

0 0 0 0 0 0 0 0 0 0 0
0 M M M M M M M M M 0
0 1 1 1 1 1 1 1 1 1 0
0 M M M M M M M M M 0
0 0 0 0 0 0 0 0 0 0 0

```

```

      0 0 0
      0 0
      0 0 0
0 0 0 0 1 0 0
0 1 1 1 1 1 0
0 0 0 M 1 M 0
      0 M 1 0 0
      0 M 1 0
      0 0 0 0
      0 0

```

```

      0 0 0 0 0
0 0 M 1 M 0 0
0 M M 1 M M 0
0 1 1 1 1 1 0
0 M M 1 M M 0
0 0 M 1 M 0 0
      0 0 0 0 0

```

(a)

```

0 0 0 0 0 0 0 0 0 0 0
0 M M M M M M M M M 0
0 E E E 1 1 1 E E E 0
0 M M M M M M M M M 0
0 0 0 0 0 0 0 0 0 0 0

```

```

      0 0 0
      0 0
      0 0 0
0 0 0 0 E 0 0
0 E E E E E 0
0 0 0 M E M 0
      0 M E 0 0
      0 M E 0
      0 0 0 0
      0 0

```

```

      0 0 0 0 0
0 0 M E M 0 0
0 M M E M M 0
0 E E E E E 0
0 M M E M M 0
0 0 M E M 0 0
      0 0 0 0 0

```

(b)

Fig. 36. End Erosion Shape Discrimination

state "M" cells specify the cells of the residue which were removed during the skeletonizing (MAT) procedure. Note that a simple cell count at this point would not permit the processor to discriminate between these shapes because they all have 9 cells in their skeletons. To improve the ability of the processor to discriminate between rectangular skeletons (skeletons which typically have a single long limb) and objects which have multiple limb skeletons, an end erosion operation (EERO) was implemented in the branches of the processor which were designed to measure rectangular features. Figure 36(b) illustrates the results of removing 3 state "1" cells from the ends of each of the skeletons of Fig. 36(a). Since the skeletons of the round and the irregular shaped objects had multiple limbs, the number of state "1" cells in each limb was small and the skeleton was completely eroded by the end erosion operation. However, since the single limb of the rectangular object was long, 3 state "1" cells remained following the EERO. Thus, the EERO has provided a convenient means of providing an improved shape discrimination capability to the cellular logic processor.

In summary, the rectangular images are measured in both width and length by applying an erosion sequence in conjunction with the medial axis transformation. An end erosion operation is then applied to improve the processor's rectangle shape discrimination capability. The next section will describe the decision criteria and how they combine with the feature extraction algorithms to support low probability of error and false alarm rate goals.

The Decision Criteria

While the earlier sections of this chapter have concentrated on

describing the algorithm selection process and the numerous innertwining relationships cellular logic feature extraction approaches exhibit, this section will describe the criteria by which decisions are made. Before proceeding, however, Fig. 37 provides a summary of the processor architecture which illustrates the 3-D noise filter, the thresholding operations, the selected erosion sequences, and the MAT and EERO operations for the three branches designed to detect rectangular objects. Following these shape measurement operations, the feature accept or reject decision is effected by comparing the residue count (r_1) of each branch to an acceptance window (AW1) designed specifically for that branch. Finally, following this comparison, the output states of all six branches (labeled L1 through L6 in Fig. 37) are logically compared to affect a target class decision.

The decision criteria proposed in the following two sections can be summarized as follows:

(a) develop accept-reject regions for each feature and target class by applying a likelihood ratio test to the class conditional pdf's experimentally obtained via introducing a synthetic training set of targets to the processor of Fig. 37.

(b) logically combining each branch's output state (feature detected or not detected) to affect a target class accept-reject criteria. This approach to designing decision criteria is often used in pattern recognition tasks because it is simple to develop and implement and also permits a great deal of flexibility in accommodating a priori information concerning the reliability of feature detection between target classes. Since the proposed decision criteria does not, for example, apply a likelihood ratio test to the joint (for all




Fig. 37. Basic Processor Design

branches/features) class conditional pdf's, the proposed design criteria performance may not be equal to that theoretically possible (Ref. 5). However, as will be described later, the flexibility attained by applying the proposed decision criteria permits adjustments to be made to the decision logic that can significantly improve the potential probability of detection without significantly increasing the error rate. The feature accept or reject decision criteria will first be described.

Feature Acceptance Window Design. As earlier described in this chapter, each branch of the six branch processor is designed to detect objects which are similar in size and shape to a feature of one of the three target classes. Following the shape measurement algorithms indicated in Fig. 37, a basic residue cell count operation is implemented. Since every object presented to the processor passes through each of the six branches, a set of three sample class conditional probability density functions (pdf's) can be estimated for the residue counts of each branch by presenting digitized randomly positioned training samples of each target class to the processor. Figure 38 provides an example of a set of class conditional pdf's for the i^{th} branch of the processor where the abscissa value indicates the residue count (i.e. the number of state "1" cells which remain) following the feature measurement algorithm and the ordinate value indicates the relative frequency of occurrence of each abscissa value experienced for 35 training samples of each target class. Note that the sample pdf's obtained in this manner are conditioned on a target class and apply only to the i^{th} branch since the threshold values and feature measurement algorithms vary from branch to branch. Since Fig. 38




Fig. 38. Sample Class Conditional Probability Density Functions (1-th Branch)

TABLE II BRANCH ACCEPTANCE WINDOWS (Ref. Fig. 37)	
AW 1 -- [1,2,3,4]	AW 4 -- [4,5,6]
AW 2 -- [2,3,4]	AW 5 -- [4,5]
AW 3 -- [1,2,3,4]	AW 6 -- [2,3,4,5]

actually represents the sample pdf's for the third branch which is designed to detect the smallest circular feature (i.e. the tops of class 3 objects), it is not surprising that the residue counts of class 2 and class 3 objects are quite similar. This is, of course, due to the fact that the tops of class 2 and 3 objects are similar in size. Likewise, since the binary images obtained when thresholding class 1 objects at the branch three threshold height are large, the branch 3 erosion sequence does not erode a large portion of the class 1 thresholded tops, and the residue counts are large and easily discerned from the eroded tops of class 2 and 3 objects. When two class conditional pdf's share common abscissa values as do classes 2 and 3 of Fig. 38, the application of a likelihood ratio test (Ref. 27) results in an acceptance window of (1,2,3,4) for the circular feature of class 3 objects. When used in conjunction with a symmetrical cost function and equal a priori class probabilities, the likelihood ratio discriminant results in minimum misclassification error rate performance which is consistent with processor design goals.

In summary, the feature acceptance window for each branch is designed by presenting 35 randomly positioned digitized training samples of each target class to the processor, estimating the class conditional probability density functions for the residue counts observed in each branch of the processor, and then applying a likelihood ratio test to establish the residue counts that are acceptable for an unknown object to be classified as containing the feature associated with each branch of the processor. Each branch of the six branch processor will thus have a single acceptance window, and an unclassified object will be determined to exhibit a specific feature only if the residue count of

that object is an element of the feature acceptance window of the branch designed to detect that feature. The acceptance windows for the processor of Fig. 37 are summarized in Table 2. If nontarget objects (clutter, noise, trees, other vehicles, etc.) are presented to the processor, the residue counts exhibited by each branch may fall anywhere on the abscissa depending upon the statistical geometries of the nontarget objects. Thus, narrow acceptance windows become highly desirable to minimize the probability of false alarms, and the earlier techniques used to determine which erosion sequences provided small residue set spreads should now be better understood. Once an object has been processed by all six branches, a determination must be made as to whether the unknown object is a member of one of the three target classes.

The Target Class Decision. Two target class decision criteria will be presented. The first requires exact classification and decreases the probability of error at the expense of detection performance to a point where a P_d of 1.0 is not possible even in the absence of noise. The second decision logic approach is less restrictive and provides for significantly improved performance by accommodating feature similarities between target classes.

Following the comparison of each branch's residue count with its acceptance window, the lines labeled L1 through L6 in Fig. 37 will be defined to be in state "T" if the residue count is an element of that branch's acceptance window and in state "F" otherwise. Thus, for an unidentified object to be classified as a member of a specific target class, the most restrictive decision criteria would require that the two branches designed to detect the top and base of that target class be in

state "T" and the remaining four branches be in state "F". For example, the processor would classify an object as a member of target class 3 if the outputs of the third and sixth branches (L3 and L6) were in state "T" and the branches L1, L2, L4, and L5 were all in state "F". Table 3 summarizes this strict target class decision criteria. While simple in design and implementation, this strict decision criteria does not take into account that the tops of class 2 and 3 targets are very similar in size as are the bases of class 1 and 2 targets. By relaxing the decision criteria to permit the target class acceptance states as described in Table 4, the performance of the six branch processor can be significantly improved. Of course, since the acceptance states of Table 3 are a subset of the acceptance states of Table 4, the P_d of the relaxed decision criteria will be higher than that of the strict decision criteria. On the other hand, since more acceptance states are allowed in the relaxed decision criteria, its false alarm rate will be larger than that of the strict decision criteria. Quantitative analysis of these relationships can, in general, only be accomplished by extensive field tests or, when clutter and noise models are available, by Monte Carlo simulations. However, one would not expect the relaxed decision criteria to exhibit a noticeable increase in misclassification rate because the tops of class 1 targets and the bases of class 3 objects are easily separable due to their unique sizes. A qualitative performance comparison of the strict and relaxed decision criteria as applied to an arbitrary measure of false alarm rate is provided in the following chapter.

In summary, the processor design approach described in this chapter provides a systematic approach to neighborhood transformation pattern

TABLE III
STRICT DECISION CRITERIA

		ACCEPTABLE BRANCH STATES					
		L1	L2	L3	L4	L5	L6
DECISION	Tgt. 1	T	F	F	T	F	F
	Tgt. 2	F	T	F	F	T	F
	Tgt. 3	F	F	T	F	F	T

TABLE IV
RELAXED DECISION CRITERIA

		ACCEPTABLE BRANCH STATES					
		L1	L2	L3	L4	L5	L6
DECISION	Tgt. 1	T	F	F	T	F	F
		T	F	F	T	T	F
		T	F	F	F	T	F
	Tgt. 2	F	T	F	F	T	F
		F	T	F	T	T	F
		F	T	F	T	F	F
		F	T	T	F	T	F
		F	F	T	F	T	F
		F	F	T	T	T	F
		F	T	T	T	T	F
		F	T	T	T	F	F
	Tgt. 3	F	F	T	F	F	T
		F	T	T	F	F	T
		F	T	F	F	F	T

recognition algorithm selection. Inherent to the systematic search design approach are several unique and productive algorithm measures of effectiveness which not only accommodate the complex geometrical probabilities associated with the application of discrete neighborhood structuring elements to digitized randomly positioned targets, but also provide a precise method of relating various performance parameters to the algorithm selection process. The concept of describing neighborhood transformations as structuring element operations has been described in the literature (Ref. 11) and is a significant aid to understanding the general effects of applying various neighborhood transformations to a pattern recognition task. However, the unique concepts of snug fitting structuring elements and residue set spread functions provide a systematic approach to quantifying algorithm optimality with respect to the desired performance criteria. The snug fit criteria insures that, in a noise free environment, an algorithm is selected that detects all targets presented to it and thus makes the potential probability of missing a target zero. The residue set spread, on the other hand, provides a method of directly relating the algorithm's potential false alarm rate to an assumed or measured model of clutter or nontarget objects. Additionally, minimizing the residue set spread is complementary to the goal of high probability of detection since it encourages selection of an algorithm which provides a consistent measure of a target object and is relatively invariant to target translation and rotation. Finally, the design approach is complemented by the use of a likelihood ratio test which provides a well understood method of relating misclassification rates to the overall processor design. The proposed design approach can be readily extended to the measurement of

arbitrary shaped geometrical objects in two or three dimensions.

This completes the description of the processor design. The next chapter will describe a Monte Carlo performance analysis of a processor designed in accordance with the procedures and design rules established in this chapter.

V. A Performance Analysis

This chapter describes the results of a Monte Carlo performance analysis of a target classification processor constructed according to the algorithm selection procedures and decision criteria established in Chapter IV. The Monte Carlo analysis is not designed to rigorously characterize the performance of the proposed processor design approach but is designed to provide insight to the sensitivities of the design and to investigate the validity of some of the assumptions made during the design process. For example, the sensitivity of the design to data which has been corrupted by additive noise or has been distorted by platform instabilities is always of interest to the algorithm designer. Likewise, an arbitrary measure of false alarm rate would provide a means of comparing the strict and relaxed decision criteria. The performance analysis also provides an opportunity to compare issues such as the relative performance of the processor in the presence of white (statistically independent) vs. correlated additive noise and also encouraged the development of a normalization technique which permitted a direct comparison of the processor's performance as applied to features which contained different signal (energy) content. Thus, while the performance analysis does not attempt to rigorously characterize the performance of the design approach of Chapter IV, it does provide much useful insight to the utility of neighborhood transformations in general as well as the specific design proposed.

Performance Defined

The performance of the processor shall be defined as the relative frequency at which the processor correctly identifies the class to which

an unlabeled target belongs. While the probability of false alarm is not accommodated in this definition of performance, a comparative measure of the false alarm rate for both the strict and relaxed decision criteria will be provided later in this chapter.

Ideally, a typical performance curve would describe probability of detection (P_d) as a function of a parameter which is independent of the specific features selected such as a ratio of the noise and signal (n/s). In such a ratio, the signal (s) selected must be appropriately related to the signal content of the features selected and the target models from which they are extracted. The noise (n) portion of this ratio is commonly a measure such as the standard deviation of the noise source, and the noise is incrementally increased from zero until the processor performs poorly. The desirability of presenting performance as a function of a noise-to-signal ratio and the concept of using only the signal level associated with the performance limiting feature can be best illustrated by the following simple example.

Given a 3-D target model and a processor which has two branches designed to extract and measure the circular and rectangular features as described in Chapter IV, Monte Carlo simulations were used to determine the set of integers (the residue counts) for each feature/branch as the target model was randomly positioned on the digitizing grid. Using these acceptance windows and the same sequence of random positions, white Gaussian noise was added to the 3-D data and the relative frequency at which the noise corrupted 3-D data exhibited a residue count which fell within the acceptance window for each feature was plotted in Fig. 39(a) for various levels of noise. Note that the circular feature was significantly more sensitive to the additive noise




Fig. 39. Application of Noise to Signal Ratio

than was the rectangular feature. This difference in sensitivity is, of course, directly related to each feature's threshold height above background since the 2-D binary image of each feature can only be distorted if the additive noise changes a background cell to exhibit a height above the threshold value or a cell which is normally (without noise) above the threshold to exhibit a height lower than the threshold. Thus, the rectangular feature has a signal level equal to one-half the height of the rectangular base of the target model since this is the rectangular feature's threshold height above background. The circular feature, however, exhibits a signal level equal to its threshold height minus the height of the top surface of the target's rectangular base. In the target model used in this example (and for the target models of Fig. 22), the signal of the rectangular feature was about three times the signal of the circular feature. Figure 39(b) illustrates that by normalizing the performance curves of Fig. 39(a) to each feature's signal level (threshold height above background) the processor can detect circular features nearly as well as it detects rectangular features. Thus, by plotting the performance curves as a function of the noise to signal ratio, a feature independent view of the processor's performance can be established. Finally, if the strict decision criteria of Table 3 is invoked, (i.e. both features of a target must be exactly detected), the ability to detect a target can be no better than the processor's ability to detect the performance limiting (smallest signal) feature. Thus, when both features are required parameters in the decision process, the processor's performance should be normalized to the signal level associated with the performance limiting feature. Normalizing the joint performance to a larger signal value would make

the results unnecessarily pessimistic and inappropriate for performance comparison with processors which incorporate different features. In summary, the signal level associated with the circular features is the performance limiting signal and the processor's overall performance curves will be normalized to the average signal content of these features.

The Analysis Procedures and Results

Monte Carlo techniques were used to both train the processor (i.e. design the acceptance windows for each branch) and to evaluate the processor's performance. The 3-D data used for both training and evaluation purposes was analytically generated by simulating a nadir viewing line scanning range sensor and incorporated the ability to simulate sensor platform roll rate, pitch rate, and altitude variations. A typical performance analysis would first generate 105 samples of 3-D data (35 randomly positioned samples of each target class). Each of these samples would be corrupted by additive noise and then be processed by the processor. The performance parameter P_d would then be estimated to be the relative frequency at which the samples were correctly classified by the processor.

Geometrical probabilities and discrete sampling constraints dictated that the algorithm training set contain rotated as well as translated target images. To accommodate rotation, the target images were sampled at seven equally spaced intervals between zero and 45 degrees to approximate a uniform distribution. Each target image was also randomly translated with respect to the digitizing grid (uniformly distributed over the interval [0,1]) to provide a representative sample of digitized geometrical images. The total sample size of 35 randomly

translated and rotated target images was the result of a compromise between computational constraints and the always desirable large sample size. Several algorithm training and evaluation sample sizes of between 10 and 100 randomly translated and rotated targets were tried and the figure of 35 was settled upon because a larger sample size did not generally change the feature's class conditional relative frequency of occurrence plot (pdf's), the feature's accept-reject regions, or the processor's overall performance. As an additional check, the random translation number generator seed was also changed and, for sample sizes of 35 and larger, the design and performance of the processor remained stable. Thus, based upon these experiment observations, the sample size of 35 randomly translated and rotated target images for each target class was established as reasonable to demonstrate the utility of the processor design approach.

The Basic Processor Performance. As defined earlier, the performance of the processor is the relative frequency at which the algorithm under evaluation correctly classifies a target object. Figures 40(a) and (b) describe the performance of the six branch target classification processor of Fig. 37 for the strict and relaxed decision criteria of Tables 3 and 4 respectively. These results are obtained by adding increasing amounts of Gaussian distributed noise to the synthetic randomly translated and rotated target samples which were used during training of the processor and then representing these distorted images to the processor. While not illustrated, the introduction of correlated noise to the target samples did not as adversely affect the performance of the processor as did the statistically independent noise. Several different random number (noise) generator seeds were used for




Fig. 40. Basic System Performance

comparison purposes and in no case did the performance of the processor vary by more than 7% from the performance illustrated. Observe that the relaxed decision criteria enabled the processor to attain a Pd of 1.0 in a noise free environment whereas the highest attained performance for the strict decision criteria is 0.75. This was due to the inability of the strict decision criteria to use a priori information concerning similarity of features within target classes 1 and 2 and classes 2 and 3. Also observe that for the features selected and the noise applied, the 3-D local average and 3-D closure noise filters both perform significantly better than the 3-D opening filter. As described in Chapter IV, this characteristic was due to the initial 3-D erosion implemented with the 3-D opening operation which removes small convex surface irregularities (such as the truncated conical tops of the targets) when corrupted with a negative noise pulse. As hypothesized in Chapter IV, the relaxed decision criteria provided for significantly improved performance. However, without a comparison of relative false alarm rates, the cost (increased false alarm rate) of implementing the relaxed criteria cannot be properly presented.

False Alarm Rate. Without validated models of clutter and nontarget objects, it is not generally possible to obtain a quantitative estimate of false alarm rates. However, the topographical contours of many geographical areas can be used to suggest modeling 3-D terrain data as a two-dimensionally correlated random process. While most terrain cannot be exactly modeled as such (Ref. 4), a two-dimensionally correlated Gaussian random process does provide a random medium against which the relative false alarm rates of the strict and relaxed decision criteria can be compared. For example, if the technique described by

Moshman (Ref. 17) is used to generate an array of correlated noise, and this noise is thresholded at the two heights, specified by the target models for the rectangular and circular features, the resulting binary image could be input to the processor and the resulting number of false alarms would provide a measure to compare the relative false alarm rates for the strict and relaxed decision criteria.

To accomplish this, 100 60x60 samples of 2-D correlated Gaussian noise with correlation distances ranging from one to five cells and standard deviations between four and sixteen cells were applied to the processor of Fig. 37. Figures 41(a) and (b) illustrate the results of applying one such noise sample to the 2nd and 5th branches respectively. The state "1" cells are the cells which remained following the feature extraction algorithms and the state "E" cells in Fig. 41(a) represent the skeleton cells which were eroded (changed to state "0") by the end erode operation. Note that in the lower central portion of Fig. 41(a) there is a binary object which eroded to a residue count of three state "1" cells and in Fig. 41 (b) the same area (at the higher threshold for the circular feature) the residue count was five. Thus, Fig. 41 illustrates one of the few samples of noise that exhibited a false alarm; a class 2 target for either the strict or relaxed decision criteria. Out of the 100 samples of correlated noise, only two exhibited false alarms for the strict decision criteria and three more false alarms were observed when the relaxed decision criteria was applied to the same samples of noise. While it is rather arbitrary to apply confidence measures to these results (since the size of the noise samples is large compared to the targets), if one considers that each sample contains four 30x30 samples of noise (and any of the targets can




Fig. 41. False Alarm Noise Samples

easily fit into this size matrix), then there are a total of 400 noise samples and the probability is at least .95 that the "true" false alarm rate for either decision criteria is less than .05 (Ref. 5). Thus, the relaxed decision criteria increased the false alarm rate but still provided a relatively small false alarm rate in the presence of two-dimensionally correlated Gaussian noise.

Additional Performance Characteristics. The previous analyses have assumed perfect control or knowledge of the dynamic characteristics of the sensor platform. The sensitivity of the proposed processor design to geometrical distortions of the 3-D data which may be induced by uncorrected sensor platform pitch rates, roll rates, or altitude errors will now be presented.

Figure 42 describes the performance of the processor of Fig. 37 (using the strict and relaxed decision criteria) when the resolution (scale) of the 3-D data is varied by as much as 5 percent. This form of distortion will occur if the sensor platform is not at the proper altitude above the terrain or if a nadir sensor is oriented slightly off vertical. The nonsymmetrical performance degradation for equal increases (+) or decreases (-) in target scale are due to the discrete nature of the available symmetrical erosion sequences. While it is obvious that the curves of Fig. 42 can be used to determine the sensitivity of the processor to altitude fluctuations of the sensor platform, they can also be used to assist in selecting a platform altitude which provides a balanced performance degradation for equal uncontrolled variations around the nominal platform altitude. Finally, while not conclusive, these curves indicate that the feature extraction algorithm selection procedure described in the previous chapter can be




Fig. 42. System Performance with Target Scale Variations

used to separate geometrical objects which differ in scale by about 2.5 percent. While a sensor platform which exhibits a constant altitude error will result in improperly scaled 3-D data, the 3-D data may also be distorted by sensor platform pitch and roll rates.

Figure 43(a) summarizes the sensitivity of the processor of Fig. 37 to several different values of roll rate using the relaxed decision criteria. The roll rates indicated are derived from several assumed parameters used in the simulation; an angular resolution of 10^{-3} radians, a dwell time (per pixel) of 10^{-6} seconds, and a row (scan line) length of 1000 pixels. Figure 43(b) provides an example of the thresholded base of a class 3 target for a stable platform and a platform exhibiting a roll rate of .5 rad/sec. Figure 44 provides similar performance characteristics for the same processor but in this case, for various levels of platform pitch rate. The increased sensitivity of the processor to pitch rate is due to the assumption that the forward motion of the sensor platform provided the scanning motion in the direction of flight. Thus, while the entire target would be scanned at full resolution under a platform roll rate environment, pitch rate would induce an undersampling of the targets. As a result, roll rate resulted in skewed objects with proper dimensions and pitch rate resulted in skewed objects with scale variations in one dimension. This effect would be observed in any line scanning data collection system.

Since the acceptability or unacceptability of the Chapter IV processor performance (as determined during this Monte Carlo simulation) is a dynamic issue and is dependent upon the appropriateness of the assumptions made during the design and the risks and costs associated with the specific operational application of the processor, it is




Fig. 43. System Performance with Uncorrected Roll Rate




Fig. 44. System Performance with Uncorrected Pitch Rate

appropriate to address these issues before proceeding to conclusions and recommendations.

Real World Considerations

The results presented in the previous sections deserve additional discussions concerning the application of such a system in a less academic environment.

The Flat Terrain Assumption. Early in this investigation, the assumption was made that the terrain was flat so that the "first order" sensitivities of the proposed processor design could be observed. While some topography may be compatible with this assumption, most geographical areas would not. This, of course, introduces a potential design deficiency since the feature extraction technique (height threshold) generally requires that the targets be placed on a surface which is normal to a vertical line connecting the sensor to the target so that the extracted features are rectangles and circles. However, even though the flat terrain assumption appears quite limiting, the existence of an imaging range sensor on board the scanning platform does provide much of the information necessary for the carrier platform to follow the contour of the terrain as accurately as its aerodynamics and guidance system permits. Thus, while this may be an unrealistic request for a large aircraft, a small lighter vehicle such as a remotely piloted vehicle or cruise missile may well be capable of adjusting its trajectory quickly enough to remain at a relatively constant altitude above some smoothly varying terrain surfaces and, in doing so, the scanning platform would remain approximately parallel to the surface. As alternatives, a random bias filter has been suggested by Sternberg (Ref. 24) which essentially performs a high pass terrain removal

operation on 3-D data, and Miller (Ref. 15) has suggested that three-dimensional features (volumes/solids) and shape measurement techniques would be less sensitive to the varying terrain heights. Note that the snug fitting algorithm selection process described in Chapter IV could easily be adapted to three-dimensional features. Thus, while the flat terrain assumption may not be representative of typical operational environments, reasonable vehicle designs, tactics, and terrain removal filters could make the proposed simplistic algorithm and processor design approach quite applicable to operational implementation.

The Nadir Assumption. The second assumption which has operational implications is that of the nadir (down looking) range measuring sensor. Historically, tacticians have preferred to look forward rather than down so that the targets are detected (and attacked or avoided) prior to the scanning platform arrival. Logically, this improves the survivability of the platform and permits the targets to be attacked head-on (like the well known Kamikaze tactic). Commercial enterprises have yet to solve the deceptively simple task of constructing a machine which can read uncontrolled printed text as well as a typical eight year old child. When this task is compared to that of classifying tactical targets which are located hundreds of meters away and are controlled and disguised by equally intelligent human beings, the difficulty of the task becomes quite apparent. Therefore, when technology has provided a data base which is relatively insensitive to diurnal variations and nature has provided a viewing position (nadir) which drastically simplifies the geometries associated with the target acquisition task, tactics, vehicles, and munitions should be developed to exploit these

opportunities. Perhaps the first truly reliable and autonomous target acquisition and classification capability could then be demonstrated.

Angular and Range Resolution. The angular resolution of the sensor, when combined with the sensor platform altitude and look angle, determines the cell size of the digitization grid. The resolution assumed for this analysis resulted in the smallest target containing about ten cells for the smallest circular feature and over two hundred cells for the larger rectangular features. The range resolution was arbitrarily set at one tenth the angular resolution since range can be very accurately estimated by laser ranging devices. Whether tactical vehicles can be afforded the luxury of such high quality sensors is a technology issue better left for others to resolve. The observation to be drawn from the performance results is that given angular and range resolutions of this magnitude, the autonomous classification of tactical vehicles appears quite feasible - - even when conservative (with respect to low probability of error and false alarm rates) processor design rules are implemented.

Feasibility of Implementation. So far, the design of a cellular logic target classification processor has been proposed and its performance has been parametrically evaluated via Monte Carlo simulation for sensitivities to additive Gaussian noise, scale, roll rate and pitch rate. This section will address the feasibility of using "off the shelf" hardware to satisfy the scale, pitch and roll rate limitations imposed by the conservatively designed processor.

The performance curves of Fig. 41 indicate that the processor can accept scale variations in the 3-D data as large as 2.5 percent without drastically reducing the performance of the processor. For a nadir

viewing sensor at an altitude of 10^3 cell diameters, this would imply that the sensor platform would have to remain within an altitude window of ± 25.0 cell diameters to insure that the scale of the 3-D data remained within the ± 2.5 percent allowable tolerance. The ability of an airborne platform to remain within such an altitude window is, of course, dependent upon the control characteristics of the pilot or autopilot as well as the weather and the nature of the terrain. Fortunately, the accurate range (altitude) information available to the pilot/autopilot should, weather permitting, make the task feasible for a reasonably responsive vehicle.

With respect to the roll and pitch rates, sensors are generally mounted on stabilized platforms within the carrying vehicle to isolate the sensors from the dynamics of the vehicle. Figures 43 and 44 illustrate that pitch rate is more restricting than roll rate and requires to be less than .05 rad/sec for reasonable processor performance. This, of course, is a very easy specification for modern day stabilized platforms to satisfy, and even suggests that, under these assumptions, a stabilized platform may not even be needed. However, more system oriented investigations will be required before such a suggestion can be seriously considered.

Acceptable Performance. The definition of "acceptable performance" in an operational environment must remain flexible due to the values and risks associated with various targets and scenarios. If an extremely low false alarm rate is required, such as when friendly forces are in the area, a probability of detection significantly less than 1.0 may be quite acceptable. Conversely, if the targets were in a free fire zone, widening the acceptance windows and/or deleting the performance limiting

feature and relaxing the decision criteria could provide a probability of detection of 1.0 for relatively large noise-to-signal ratios. Of course, a high P_d obtained in this manner may well result in the destruction of a large number of trees, rocks and barns so such a weapon must be cheap and plentiful. Thus, the decision was made in Chapter IV to design the processor to reasonably low false alarm and error rates, and to require the exact separation of the three target classes. This provides a pessimistic estimate of processor performance but leaves the option to adjust various parameters to increase the probability of detection or probability of false alarm rates as tactical requirements allow.

Discussion of Results

The results of the Monte Carlo performance analysis confirm the utility of developing target classification algorithms in accordance with the "cookbook" procedure described in Chapter IV. In general, the performance of the design was severely degraded when additive Gaussian noise was added to the 3-D data so that the noise to signal ratio was greater than 1.0. The performance was also significantly influenced by the use of 3-D noise filters. While it is clear that the 3-D opening noise filter is not suitable for use with the features selected and the noise as modeled, either the 3-D local average or the 3-D closure noise filters appear to perform quite well, but additional testing would be required to establish which of the two is best. Likewise, the relative insensitivity of the proposed design to platform instabilities of roll and pitch rate are very promising, but the absolute extent of these trends remain unresolved. However, it is clear that a relatively inexpensive stabilized platform may be all that is needed to isolate the

sensor (data) from platform motion. On the other hand, scale (altitude) variations of more than $\pm 2.5\%$ significantly reduced the performance of the processor. Sensitivity to scale variations are common to most pattern recognition tasks. If additional insensitivity to scale variations is required, scale invariant features such as length-to-width ratios or invariant moments could be used. The general algorithm development approach of Chapter IV could readily be adapted to these features as well. Finally, the arbitrary measure of false alarm rate was not designed to establish performance as such, but was included to measure the relative difference between the strict and relaxed decision criteria.

Since a generally applicable definition of acceptable performance does not exist, the performance of an algorithm designed in accordance with Chapter IV procedures must be carefully evaluated with respect to the task at hand. Specifically, one must compare the risks associated with not detecting a target to the costs of declaring a nontarget object a target of interest. Since such an evaluation was not intended, and, recognizing that costs and risks are often variable or difficult to define, the Chapter IV design approach has provided an algorithm selection technique which provides a relatively low false alarm and error rate. Thus, if improved performance is desired, a relaxed decision criteria can be applied or the feature acceptance windows can be widened. Of course, any improved performance attained in this manner would be at the risk of increased potential false alarm rate. As indicated at the beginning of the chapter, the purpose of this performance analysis was not to exhaustively characterize the performance of the processor but was designed to demonstrate the

potential utility of the Chapter IV neighborhood transformation algorithm design approach and to characterize its relative sensitivities to several common forms of noise and data distortions.

VI. Conclusions and Recommendations

Conclusions

The following conclusions can be made concerning the uniqueness of three-dimensional range data, the snug fitting feature measurement algorithm selection procedure, and the results of the Monte Carlo performance analysis.

1. Actively collected high resolution range data provides a significant technical opportunity for the designers of tactical target classification seekers because:

a. The data is unique in that it is relatively free from the diurnal variations which have historically complicated the tasks of noise removal and feature selection and extraction in passively collected image data.

b. The multi-level information contained in 3-D range data is directly related to the three-dimensional shapes of objects within the scene rather than the thermal emissivity or optical reflectivity of the scene. This permits the use of features which are direct measures of the three-dimensional shapes of the targets.

c. The consistency of the data and the data's direct measurement of geometrical shapes provides for the realization of simple, well understood, and believable target, nontarget, and background/clutter analytical models which may be used in extensive Monte Carlo simulations for performance analysis. The potential for such a realizable target/clutter model will permit analytical estimates of a system's performance,

including false alarm rates, to be estimated without initiating extensive flight tests for data collection purposes.

2. The use of two- and three-dimensional equivalent structuring elements to describe cellular logic operations provides useful insight to the geometrical interpretation of sequences of erosion and dilation neighborhood transformations.

3. The Chapter IV algorithm design approach uses the unique concepts of residue set and residue set spread to permit the evaluation of a neighborhood transformation's feature extraction effectiveness to proceed in a "cookbook" manner. Thus, given an arbitrary feature extraction technique, a look up table can be developed to identify the specific sequence of neighborhood transformations which demonstrate

- a. the ability to consistently extract that feature in a noise free environment
- b. the lowest potential false alarm rate (assuming uniformly distributed clutter/false alarms), and
- c. the most relative invariance to object rotation and translation.

4. The utility of applying the cellular logic algorithm selection procedure and decision criteria of Chapter IV was demonstrated by the development of a 3-D target classification processor which, under various scenario and processor architecture assumptions, exhibits the following characteristics:

- a. the performance of the processor does not degrade significantly if the scale of the data (altitude of the sensor

platform) remains within ± 2.5 per cent of nominal and the roll and pitch rates of the sensor are kept at rates easily attainable by commercially available stabilized platforms,

b. compared to the strict application of a likelihood ratio test, the relaxed decision criteria (which accounted for the similarity of features between target classes) improved the processor's probability of detection (P_d) from .7 to 1.0 without noise and, in the presence of white additive Gaussian noise, improved the P_d from .4 to .7 at a noise-to-signal ratio of about 1.0,

c. the probability is .95 that the false alarm rate for both the strict and relaxed decision criteria was less than 5 per cent in the presence of two-dimensionally correlated Gaussian noise.

Recommendations

1. The algorithm efficiency measures of residue set and residue set spread have demonstrated the feasibility of characterizing the performance of neighborhood transformations. Additional investigations into the performance characteristics of neighborhood transformations may lead to a more powerful means of describing their performance. It is entirely feasible for graphical characterizations, similar to a receiver operating characteristic, to be developed for neighborhood transformations.

2. While the Chapter IV design approach provides a well organized means of selecting a specific sequence of neighborhood transformations to accomplish a task, it is conditioned upon a general approach (to feature measurement) being defined.

C

Unfortunately this general approach is presently developed in an ad hoc manner and, once selected, the implementation of an algorithm may not be optimal with respect to accomplishing the algorithm in the fewest number of steps/stages. Likewise, the lack of a suitable probability space for a large class of neighborhood transformations limits our understanding of algorithm performance to Monte Carlo or trial and error analysis. The development and characterization of a robust neighborhood transformation algebra would significantly improve our ability to address these issues.

3. A statistically formulated generic 3-D data model which is suitable for use in Monte Carlo simulations and incorporates terrain, clutter, vegetation, nontarget and target objects should be developed so that improved estimates of performance including false alarm rates can be established. This model should incorporate a 3-D coordinate transformation capability to provide for other than nadir look angles.

4. A processor of similar design but which extracts, measures, and accepts or rejects features entirely in three space should be evaluated for comparison purposes. This processor would use all the shape information available in the 3-D data rather than only two slices as in the Chapter IV processor design and would be less sensitive to the flat terrain assumption.

5. Simplified processor architectures which incorporate fewer than one branch per feature-target class should be evaluated for performance. While such simplifications will generally provide for less optimal algorithm selection and an increase in false alarm rate, they should be considered for processor size, weight, and

O

cost reduction purposes--especially since a generic 3-D data model, as suggested above, could be used to evaluate the P_d vs P_{fa} tradeoffs associated with these designs prior to constructing hardware.

Bibliography

1. Blum, H. "A Transformation for Extracting New Descriptors of Shape." Models for the Perception of Speech and Visual Form. Cambridge, Ma: MIT press, 1967. W. Wathen-Dunn, Ed.
2. Bullock, Bruce. The Performance of Edge Operators on Images with Texture. Hughes Research Laboratories, Malibu, Ca. AD/A-006 230.
3. Burks, A. W., Ed. Essays on Cellular Automata. Urbana, Il: University of Illinois Press, 1970.
4. Cannon, Mark W., Jr. Terrain Contour Matching (TERCOM) Sensitivity to Heading and Ground-Speed Errors. Wright Patterson AFB, Ohio: AMRL TR -77-84, May 1978.
5. Duda, Richard O. and Peter E. Hart. Pattern Classification and Scene Analysis. New York: Wiley, 1973.
6. Duff, M. J. B. "Cellular Logic and its Significance in Pattern Recognition." AGARD Proc. No. 94 on Artificial Intelligence, 25: 1-13 (1971)
7. Golay, Marcel J. E. "Hexagonal Parallel Pattern Transformations." IEEE Trans. on Computers, Vol. C-18, No. 8:733-740 (Aug. 1969).
8. Grey, Stephen B. "Local Properties of Binary Images in Two-Dimensions." IEEE Trans. on Computers, Vol. C-20, No. 5:551-561 (May 1971).
9. Harding, E. F. and D. G. Kendall, Eds. Stochastic Geometry. New York: Wiley, 1974.
10. Ingram, Marylou and Kendall Preston, Jr. "Automatic Analysis of Blood Cells." Scientific American. November, 1970.
11. Journal of Microscopy. Vol. 95, Part 2. Oxford, England. Apr. 1972
12. Kendall, M. G. and P. A. Moran. Geometrical Probabilities. Hafner, 1963. ISBN 0-02-847730-8.
13. Maruoka, A. "Cellular Automata." J. Inst. Electron. Commun. Eng. Japan: Vol. 61, No. 10:1073-1083, 1978.
14. McCormick, B. H. "The Illinois Pattern Recognition Computer - ILLIAC III." IEEE Trans. on Computers. Vol. EC-12:791-813. Dec., 1963.
15. Miller, Peter E. An Investigation of Boolean Image Neighborhood Transformations. (PhD Dissertation) Ohio State University, 1978.

16. Minsky, M. L. and S. Papert. Perceptrons, Introduction to Computational Geometry. Cambridge, Ma.: MIT Press, 1970.
17. Mashman, J. "Random Number Generation." Mathematical Methods for Digital Computers. Ralston and Wilf, Eds. New York: Wiley, 1960.
18. Nishio, H. "A Classified Bibliography on Cellular Automata Theory - With Focus on Recent Japanese References." Proc. IEEE Int. Symp. Uniformly Structured Automata and Logic: 205-214 (1975).
19. Papoulis, Athanasios. Probability, Random Variables, and Stochastic Processes. New York: McGraw-Hill, 1965.
20. Preston, K., Jr. et al. "Basics of Cellular Logic with Some Applications in Medical Image Processing." Proc. IEEE, 65-5: 826-856 (May, 1979).
21. Preston, P., Jr. and Phillip E. Norgren. "Interactive Image Processor Speeds Pattern Recognition by Computer." Electronics: October 23, 1972.
22. Rosenfeld, Azriel and C. R. Dyer. "Cellular Pyramids for Image Analysis." Department of Computer Science, University of Maryland College Park, TR544, TR596.
23. Steinberg, Richard A. "Infrared Surveillance. Part 2: Continuous-Time Signal Processors." Applied Optics, 19-10:1673-1687 (15 May 1980).
24. Sternberg, S. Unpublished communication. Environmental Inst. of Michigan (ERIM).
25. Unger, S. H. "A Computer Oriented Toward Spatial Problems." Proc. of the IRE, 46:1744-1758 (Oct. 1958).
26. von Neumann, J. and A. W. Burks, Eds. Theory of Self Reproducing Automata. Urbana, IL: Univ. of Illinois Press, 1966.
27. Van Trees, Harry L., Detection, Estimation, and Modulation Theory, Part 1. New York: Wiley, 1968.

Vita

Lawrence Arthur Ankeney was born in Pittsburgh, Pennsylvania on 1 December 1943, the son of Robert K. and Vera S. Ankeney. He graduated from North Hills High School, Pittsburgh, Pennsylvania, in June 1961. He received his Bachelor of Science from Grove City College, Grove City, Pennsylvania, in June 1965, after which he was commissioned a Second Lieutenant in the U. S. Air Force. He served as a Deputy Missile Combat Crew Commander in the 321st Strategic Missile Wing, Grand Forks Air Force Base, North Dakota until his assignment to the Air Force Institute of Technology in June 1969 where he received his Masters of Science in Electrical Engineering, Guidance and Control. He then served as a Project Engineer from 1971 to 1974 for the Department of Defense and as a Staff Engineer in the Future Systems Division, HQ Strategic Air Command, Omaha, Nebraska until his assignment to the Air Force Institute of Technology in June 1976. He is presently serving within the Guided Weapons Division of the Air Force Armament Laboratory, Eglin AFB, Florida. He is married to the former Marilyn E. Good of Pittsburgh, Pennsylvania, and they have one daughter, Emily.

Permanent address: 103 Hillcrest Drive

Pittsburgh, Pennsylvania 15237

**DATA
FILM**